Title
An Investigation in the Use of Inductive Loop Signatures for Vehicle Classification

Author
Sun, Carlos

Publication Date
2000-03-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

An Investigation in the Use of Inductive Loop Signatures for Vehicle Classification

Carlos Sun
California PATH

California PATH Research Report
UCB-ITS-PRR-2000-4

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

Report for MOU 376

March 2000

ISSN 1055-1425

AN INVESTIGATION IN THE USE OF INDUCTIVE LOOP SIGNATURES FOR VEHICLE CLASSIFICATION

FINAL REPORT

Carlos Sun

California PATH

Prepared for

Partners for Advanced Transit and Highways

PATH MOU 376

August 1999

ACKNOWLEDGEMENT

The author gratefully acknowledges the assistance provided by our research partners Drs. Stephen Ritchie and Wilfred Recker from the University of California at Irvine, Mr. Joe Palen from the California Department of Transportation, Gardner Transportation Systems, 3M Intelligent Transportation Systems, Peek Traffic Sarasota, and Mr. Rick Nelson from the City of Irvine.

ABSTRACT

This final report describes an advanced traffic surveillance technique based on pattern recognition and the use of current inductive loop technology. The focus of the investigation was a study of the feasibility of using inductive loop signatures for obtaining vehicle classification information on a network-wide level. The potential benefits from the vehicle classification information include improvements in vehicle reidentification algorithms, roadway maintenance, vehicle emissions management, roadway design, traffic modeling and simulation, traffic safety, and automatic toll collection. The compilation of a vehicle classification database is also a valuable resource for researchers in the areas of transportation planning and control. This effort complements current PATH (Partners for Advanced Transit and Highways) research in advanced surveillance technology. Different pattern recognition techniques for vehicle classification such as classical decision theoretic approach and advanced neural networks were employed in this research. Inductive signatures from the SR-24 freeway were used to test vehicle classification algorithms. Classification rates of greater than 80% were obtained using different datasets. The results demonstrated the feasibility and the potential of using this method for collecting vehicle classification data.

Keywords: automatic vehicle classification, vehicle detectors, traffic surveillance, neural networks, advanced traffic management system

EXECUTIVE SUMMARY

Vehicle classification is the process of separating vehicles according to different predefined classes. Vehicle classification information can be used in many transportation applications including road maintenance, vehicle reidentification, emissions/pollution estimation, traffic modeling and simulation, traffic safety, and toll setting. The collection of vehicle classification data has previously been investigated using various new surveillance technologies such as video, laser, and acoustic detectors. One goal of this project is to exploit the current inductive loop infrastructure. However, this fact does not abolish the possibility of changing the vehicle classification algorithms developed in this research to accept other forms of detector output in the future. In fact, detector fusion can the wave of the future for traffic surveillance once the economics of multi-detector surveillance become more reasonable.

This final report presents two distinct vehicle classification methods using the following seven vehicle classes: car, SUV/pickup, van, limousine, bus, two axle trucks, and trucks with greater than two axles. These seven vehicle classes are interesting to study because they cannot be determined by using axle counters, and they can be used in different transportation applications. Both classification methods utilize vehicle inductive signatures collected from existing loop detector infrastructure. One method uses a heuristic discriminant algorithm for classification and multi-objective optimization for training the heuristic algorithm. Feature vectors obtained by processing inductive signatures are used as inputs into the classification algorithm. Three different heuristic

algorithms were developed and yielded encouraging results of 81%-91% overall classification rates. The second methoid uses Self-Organizing Feature Maps (SOFM) for classifying vehicles. A SOFM is an Artificial Neural Network (ANN) that forms clusters of neurons which reflect similarities in the input vector. The inputs to the SOFM are the processed inductive signatures. The SOFM produced similar results of over 80% overall classification rates. These results demonstrate the potential of collecting network-wide vehicle classification data from inductive loops. The availability of vehicle classification data helps to improve traffic surveillance and paints a more detailed picture of dynamic traffic networks.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

CHAPTER 1. INTRODUCTION

1.1 OVERVIEW MOTIVATION

As new exciting technologies are being developed for transportation applications and new Intelligent Transportation Systems (ITS) are being tested and implemented, there is a growing need for more detailed and more accurate transportation data. Researchers are exploring the use of various new forms of detectors and surveillance systems. These detectors include infrared, microwave, ultrasonic, acoustic, Automatic Vehicle Identification (AVI), piezo-electric, magnetic, Global Positioning System (GPS)/cellular modem, and video image processing. Apart from investigating new technologies, there is the opportunity to squeeze out more information from the current inductive loop infrastructure. One way this is accomplished is by using inductive signatures collected from existing loops. Previously, inductive loop detector cards only operated in a pulse or a presence mode, producing a digital output. This digital output is a step function that has a duration of a vehicle's loop occupancy time or a short pulse of between 100 and 150 ms actuated by a vehicle call (ITE, 1990). However, detector manufacturers are increasingly producing detectors that have the capability to output a vehicle inductive signature through a serial port on the detector card. This inductive signature is the result of the net decrease in detector inductance when a vehicle's metallic mass passes over the magnetic field generated by the inductive loop. This signature is also referred to as the signature or the footprint of a vehicle. Some researchers distinguish this signature output from the standard pulse or presence outputs by calling it analog output. By using signatures, important transportation data such as travel time, density, partial dynamic origin/destination demand, lane change, and vehicle classification data can be obtained.

Vehicle classification is the process of sorting vehicles according to predefined vehicle classes. The Federal Highway Administration (FHWA), California Department of Transportation (Caltrans), county transportation authorities, and even private organizations such as tollway agencies all use vehicle classification in diverse ways. Therefore the set of vehicle types is defined differently for each agency or organization. Research in vehicle classification will assist public and private agencies in different areas including road management, road maintenance, data collection, multi-modal transport, emissions control, bus preemption, and toll assessment functions. Thus vehicle classification data is a valuable form of transportation data that can be used in many areas of transportation.

One particular example of a vehicle class that is useful to be monitored is trucks. Trucks and other oversized vehicles have distinctly different performance characteristics than passenger vehicles. Trucks on the average travel at slower speeds, occupy more road space, have longer maneuvering times, have longer braking distances and times, and are sometimes lane restricted. In terms of traffic flow considerations, an accurate measurement of trucks on the road will lead to more accurate modeling and simulation of real world conditions. The assumption of uniformity of vehicle characteristics is a contributor to unrealistic modeling of real networks. The ability to convey and predict traffic conditions accurately on the roadway will also lead to improved efficiency through traffic control strategies. Because the speed differential between trucks and cars is usually significant, and also because trucks are much larger than vehicles, having vehicle

classification data can help safety research and implementation. Heavier vehicles like trucks also contribute disproportionally to the wear and tear of pavements. Knowing the actual number of trucks traveled on a roadway can help estimate the life of the current road surface and assist in the scheduling of maintenance.

Another example of a vehicle class that is useful to be detected are light vehicles (LTV's). Some issues involving LTV's include the safety characteristic of such vehicles and car following behavior resulting from a higher driver field-of-view. The National Highway Traffic Safety Administration (NHTSA) reports that rollover crashes are one of the most significant safety problems for all classes of LTV's especially light trucks: pickups, sport utility vehicles, and vans (Garrott et al., 1999). NHTSA reports that rollovers resulted in an average of approximately 9000 fatalities per year or about thirty percent of all LTV fatalities from 1992 through 1996. Even though LTV's are in 68 percent as many crashes per registered vehicle as are passenger cars, LTV's are in 127 percent as many rollover crashes per registered vehicles.

Examples of applications with other vehicles classes also abound. Traffic signals can be pre-empted or extended when a bus or an emergency vehicle is detected. Automatic vehicle classification can be used for setting fees on toll roads. Vehicle class information can also help traffic enforcement agencies in maximizing their resources. Obtaining an area-wide assessment of traffic vehicle classes can help agencies estimate more accurately the types and amounts of pollution emitted by vehicles. Vehicle class information also allows the analysis of the effects of the increase in certain vehicle

classes such as the growing percentage of sports utility vehicles (SUV) in the United States.

Previous research in this field include inductive loop, video, infrared, and acoustic classification systems. There are also commercial classification systems such as the IVHS2000 by Intersection Development Corporation. The existence of commercial products should not undermine this research effort. One problem with commercial systems is the fact that detailed performance analysis and methodology is not published. The development of a PATH system will allow agencies the ability to understand the theory behind the operation of such systems, and will allow agencies to even develop their own field software according to their needs.

One advantage of the proposed system is the utilization of current infrastructure. Since inductive loops are so prevalent among both city and state roadways, there exists the possibility of a network-wide implementation of the system. Current commercial vehicle classification systems require the use of a multi-beam light curtain and a Doppler radar, or treadles and loops. Comparing with existing commercial implementations, the proposed system is much simpler with respect to equipment needs.

There is strong evidence that a given vehicle will have similar inductive signatures at different sites. A site is composed of the roadway geometry, physical inductive loop detector, lead-in cables, and detector circuitry. Even the signature changes with vehicle speed, vehicle offset (both horizontally and vertically), loop type, and other factors, there

are some changes that can be eliminated. For example, the signature can be transformed to be speed-invariant by measuring speed (double loops, radar, single loop estimate). Also for example, the nominal site inductance level, which is a function of the loop geometry and detector settings, can be normalized for preserving the shape of vehicle signatures. In addition, the nominal site inductance level can also be calibrated for each individual site so that the total amount of signal (power) can be meaningful. Since the shapes of the vehicles are preserved to some extent by different transformation and normalization methods, a value proportional to the total power is the maximum inductance recorded. This maximum inductance is a easily computable value that can be used as a feature vector. More complex Fourier or Wavelet analysis is also used for deriving the power spectrum. One evidence for the assumption that vehicle signatures at different classification sites can be similar is the work in vehicle reidentification using inductive loop detectors (Sun et al., 1998 and 1999). The vehicle reidentification results show that individual vehicles can be reidentified from one site to another site with accuracy of between 60% to 75%. In order for this to be possible, the signature features have to be preserved from one site to another.

1.2 STRUCTURE OF REPORT

The organization of this final report is as follows. Chapter 1 is an introduction of vehicle classification. Chapter 2 discusses the background and reviews preceding work. Chapter 3 discusses the data collection process and describes the database of vehicle signatures. Chapter 4 presents the theoretic decision approach and heuristic vehicle classification algorithms. Chapter 5 describes the Self-Organizing Feature Map (SOFM), which is an

Artificial Neural Network (ANN) employed for vehicle classification. Chapter 6 discusses the related issue of single loop speed estimation scheme. The use of this estimation scheme allows the extension of the vehicle classification to single inductive loops. Chapter 7 concludes this report with a summary and a discussion of future research directions.

CHAPTER 2. BACKGROUND AND PRECEDING WORK

There has been fertile research in the analysis and use of inductive signatures for transportation data collection both in the United States and abroad. Böhnke and Pfannerstill (1986) discussed the use of inductive signatures and Karhunen-Loeve transformation in reidentifying vehicle sequences. Kühne in cooperation with other researchers has published several reports of section measure instrumentation for obtaining inductive loop vehicle signatures. A freeway control system using a dynamic traffic flow model and vehicle reidentification technique was published by Kühne (1991). Another system discussed by Kühne and Immes (1993) detailed a freeway control system using section-related traffic variable detection. Pursula and Pikkarainen (1994) used a seven vehicle class scheme for classification using double inductive loop signatures and self-organizing feature map. Sun et al. (1998, 1999) describe a travel time and density measurement system using inductive loop signatures.

The following are some examples of some existing classification schemes and methodologies used by researchers and public agencies. These examples show that vastly different classification schemes exist among various transportation applications. One well-known classification scheme is the one specified by the Federal Highway Administration and is shown in Table 1 (USDOT, 1997).

Table 1. FHWA Classification Scheme

Category	Description
1	motorcycles
2	passenger cars
3	other 2 axle 4 tire singles
4	buses
5	2 axle 6 tire singles
6	3 axle singles
7	≥ 4 axle singles
8	≤ 4 trailers
9	5 axle trailers
10	≥ 6 axle trailers
11	≤ multi-trailers
12	6 axle muti-trailers
13	≥ 7 multi-trailers

Pursula and Pikkarainen (1994) used the seven vehicle class scheme in Table 2 for classification using double inductive loop signatures. Both loop signatures were used in an array configuration that improved classification accuracy over single loop configuration. The physical inductive loops used were 3x2 meters. A 12x12 Self-Organizing Feature Map (SOFM) was used with analog loop signals (inductive signatures) as inputs. The reported percentage of success was around 80% for the training set. An improved algorithm using length windows, SOFM, and Learning Vector Quantization (LVQ) resulted in a 96.2% correct classification rate with a test set size of 737 vehicles. However, only the SOFM procedure was discussed in the paper but not the LVQ procedure.

Table 2. Seven Vehicle Class Scheme

Category	Description
1	Car or van
2	Truck
3	Bus
4	Truck with semi-trailer
5	Truck with trailer
6	Car with trailer
7	Car with mobile home

Wei et al. (1996) used the three vehicle class scheme in Table 3 by applying video classification. The feature extraction process involves ground segmentation and background subtraction. The resulting image was quantized by using 16 windows and three characteristics: occupation rate, horizontal image line, and vertical image line. A backpropagation Artificial Neural Network (ANN) with one hidden layer was used. The resulting recognition rates were 98.5%, 96.92%, and 91.94% for heavy vehicles, small cars, and motorcycles using a test set of two hundred images.

Table 3. Three Vehicle Class Scheme

Category	Description
1	heavy vehicles
2	small cars
3	motorcycles

Yuan et al. (1994) used the six vehicle class scheme in Table 4 by also applying video classification. Four vehicle extraction models were used to obtain vehicle features. They are: perspective projection, length measurement, width and height estimation, and profile characters. A two-level hierarchical classification scheme was used. Level one separated vehicles into three categories by using lengths and heights. This level used the k nearest-

neighbor rule. The reported classification rate for level one was over 90%. Level two further divided vehicles by using profile characters and heuristics.

Table 4. Six Vehicle Class Scheme

Category	Description
1	2 axle vehicles: car, van, pickup, ambulance, and single-unit truck
2	buses
3	3 axle truck
4	one-unit truck, four or more axles
5	two-unit truck, four or more axles
6	three-unit truck, four or more axles

Nooralahiyan et al. (1997) used acoustic signatures to derive the four vehicle classification scheme listed in Table 5. Linear Predictive Coding (LPC) was utilized to extract feature vectors through parameter conversion from autocorrelation analysis. A Time Delay Neural Network (TDNN) was used for classification. The reported classification accuracy was 82.4% for 400 test patterns.

Table 5. Four Vehicle Classification Scheme

Category	Description
1	buses or lorries
2	small or large saloons
3	Motorcycles
4	light goods vehicles or vans

Lu et al. (1989) used an infrared detector with the four vehicle class scheme in Table 6. This paper also included the possibility of user-defined vehicle categories. A modified Sobel's edge-sharpening method was used to accentuate (or brighten) the engine, emission pipe, and wheels. Features such as size, height, and location of bright clusters,

are found through the analysis of the size and position of the bright clusters. Similar to Yuen et al. (1994), a k nearest neighbor method was used for classification into four groups. The reported accuracy was 95% for a test sample of 100 vehicles.

Table 6. Four Vehicle Classification Scheme

Category	Description
1	passenger cars
2	buses with new engines and trucks without upright pipes
3	trucks with upright pipes and trailer trucks
4	unidentified

Leonard (1991) used a ten truck category classification scheme for analysis of truck crashes including straight trucks, tractor-semitrailer combinations, and double combinations. The graphical examples for truck configuration and labeling are given in Figure 1. When developing a discriminant analysis scheme; however, the classification uses only the vehicle length as a feature vector and is performed using two categories only. For example, a vehicle is classified as a large truck or not as a large truck. Using vehicle length cutpoints of 36.1’ (11m), 39.6’ (12m), and 42.8’ (13m) the reported classification percentages were 94.2%, 96.5%, and 98.2%.

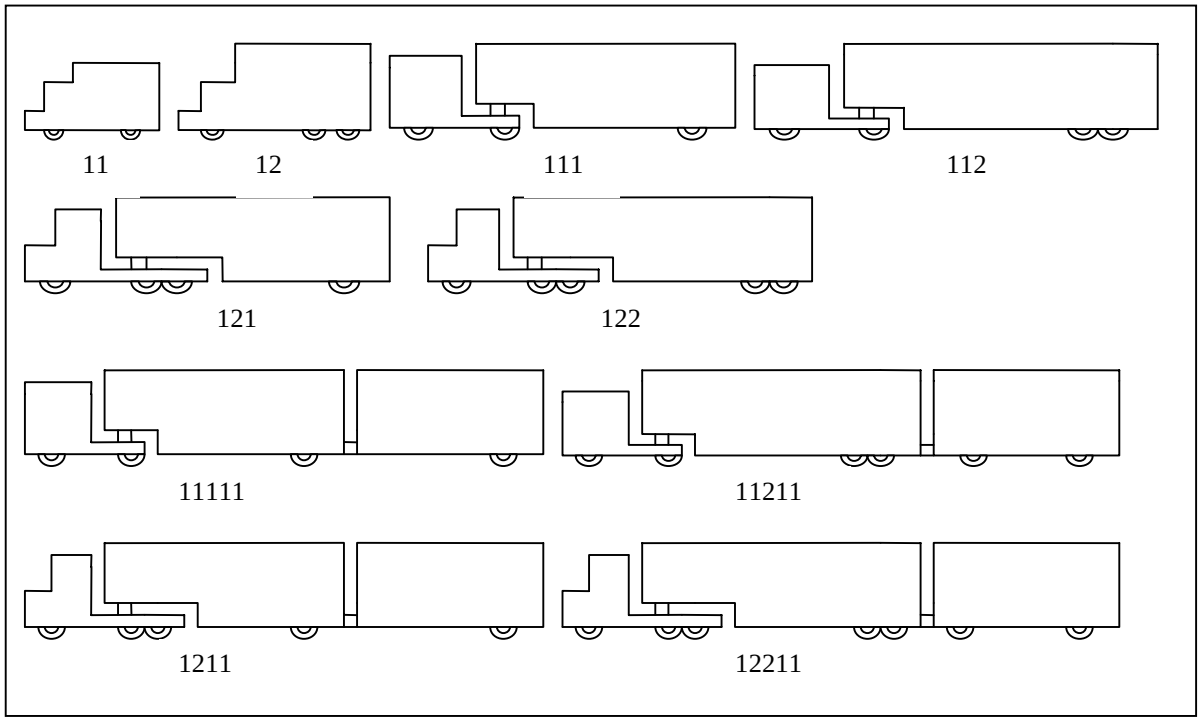

Figure 1. Truck Classification Scheme

Commercial businesses in the areas of trucking, freight forwarding, and shipping often have their own fleet classifications. Figures 2-8 shows an example of a seven vehicle classification of a commercial fleet in the United States. The classes differ in pricing, cargo capacity, and cargo versatility. Table 7 shows some class characteristics.

Table 7. Commercial Fleet Characteristics

Class	Cubic Feet	Interior Dimensions	Payload
Pickup	N/A	8’x4’ (truck bed)	1565lbs. min.
Cargo Van	305	N/A	2000lbs.
15’ Cutaway	305	N/A	2000lbs.
15’ Box Truck	822	1’-3”x7’x7’-10”	3000lbs.
18’ Box Truck	N/A	N/A	N/A
24’ Box Truck	1520	2’x8’x7’-10”	10,000lbs. min.
Stakebed	N/A	N/A	N/A

Figure 2. Pickup

Figure 3. Cargo Van

Figure 4. 15’ Cutaway

Figure 5. 15’ Box Truck

Figure 6. 18’ Box Truck

Figure 7. 24’ Box Truck

Figure 8. Stakebed

There are also commercial products available that use infrared detector or multi-beam vertical light curtain for vehicle classification. Companies such as Schwartz Electro-Optics (1998) and Transport Data Systems (1998) use 4 and 5 category schemes with subcategories. Earlier classification systems often employed a treadle, therefore the classification was based on the number of vehicle axles.

CHAPTER 3. SIGNATURE/VIDEO DATA AND DATA PROCESSING

A discussion on the signature/video data and data processing is helpful in two respects. First, the details of requirement for field implementation of the vehicle classification algorithms are beneficial to agencies who seek to deploy such a system. One objective of this research is to enable government agencies to utilize the proposed technology in the near future. Second, the careful explanation of data collection and processing procedures is provided for the use of other researchers who are interested in research in vehicle classification.

3.1 DATA COLLECTION

The data collection effort for vehicle classification is a superset of the field implementation in the sense that in addition to collecting the vehicle signature data, there is the additional need to collect video ground-truthing information. The ground-truthing is needed for researchers to evaluate the performance of their vehicle classification systems. This chapter will therefore center on the development of the data collection setup, since the field implementation is similar to data collection.

Figure 9 shows the freeway data collection site located in Lafayette, California. This site is composed of an upstream and a downstream station. Having two stations instead of a single station was necessitated by another PATH project in vehicle reidentification. Having two stations also is useful in this PATH project since it can be used to test the robustness of the vehicle classification system. The upstream loop station is located near

the Central Laayette on-ramp while the downstream loop station is before the Acalanes Road off-ramp.

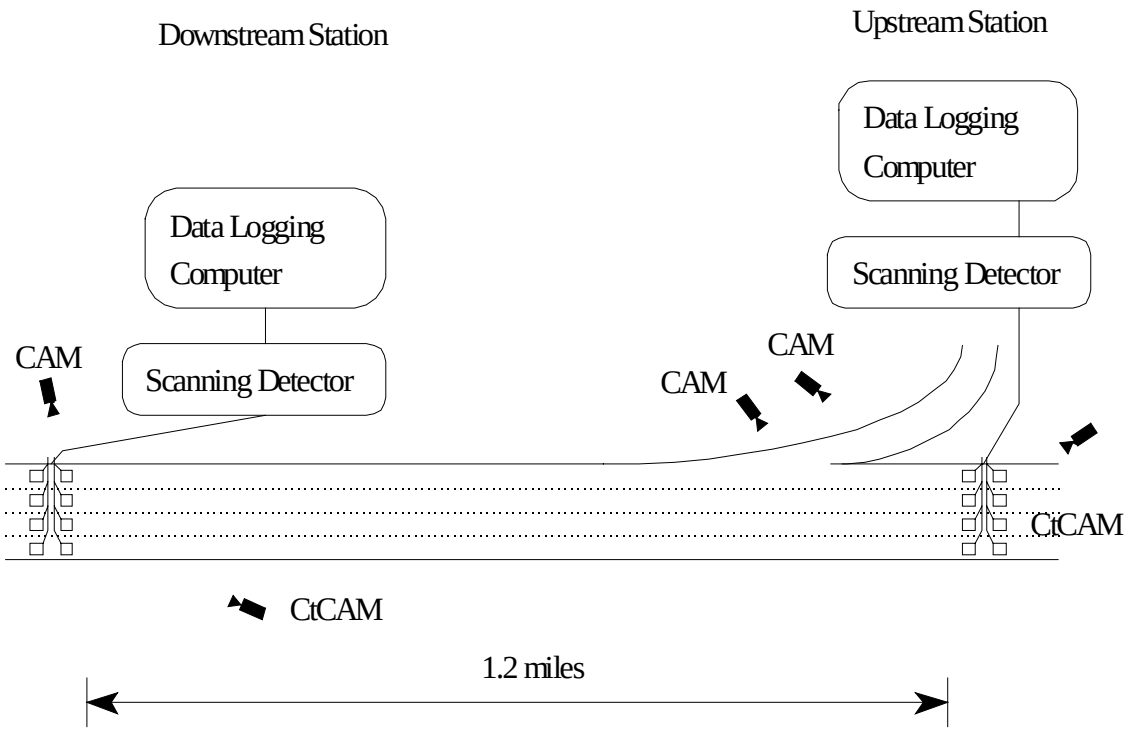

Figure 9. Freeway Data Collection Site.

The signature data were obtained from a field site on the westbound SR-24 freeway in Lafayatte, California in December 1996. Two data acquisition stations were instrumented with video, loop signature dataloggers, and speed trap dataloggers. Standard 6ftx6ft (1.82mx1.82m) loops were used at both stations. The data set contains the signatures of the upstream and downstream vehicles along with their speeds (double loop), electrical length (derived from occupancy time), arrival time at stations, and the ground truth vehicle identification number. The data set was composed of moderate flow traffic (1000 VPHPL) and contained approximately 2000 vehicles, many of which were

passenger vehicles. This data set was further reduced to 300 vehicles. The reduction was performed in order to avoid the problem of a single vehicle class (passenger car) dominating neural network weights or algorithm parameters. The reduced dataset contained a more uniform distribution of signatures across all vehicle classes. This data set was recorded on December 6 at approximately 12:00pm. The data sets were divided into training and testing data sets. The calibration of the heuristic algorithms was performed on the training data only, and the training of the Self-Organizing Feature Map was performed using a even smaller data set of 26 vehicles.

The hardware interconnect of each data acquisition station is composed of four different components: the data processing computer, time synchronization source, physical loops, and the detector cards. Figure 10 illustrates the partial hardware interconnect for instrumenting two lanes of a freeway out of four total lanes. RS-232 serial communications was used to connect various components.

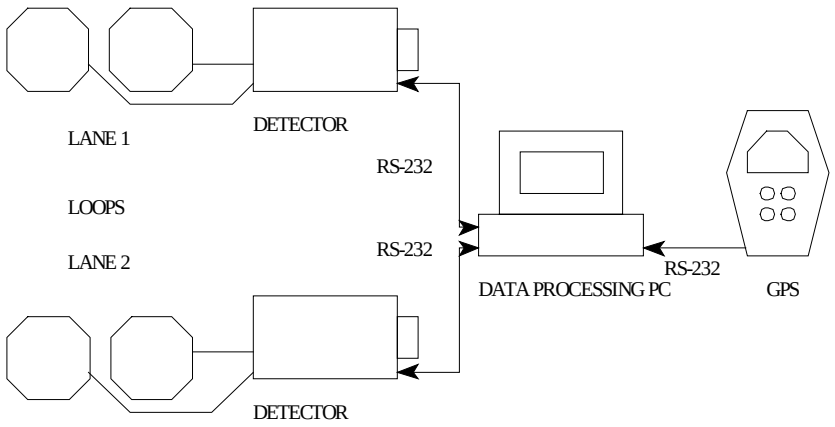

Figure 10. Data Acquisition Setup

The following paragraphs describe the signature and ground-truth video data. The signature data were stored in SIG files. The unprocessed signature files and PVR (Per Vehicle Record) files were collected in the field. The video correlation database was produced in the laboratory from the video footage taken during data collection. Manual correlation of the downstream with the upstream vehicles using video was required for the development of the video correlation database. Due to discrepancies in time stamping the loop data and the video, another manual correlation was required to match the video correlation vehicle records with the vehicle loop signatures. The last step in the process of creating the signature data file was to append unique identification numbers to the signature data that could be referenced back to the video database.

A sample vehicle signature and detailed explanation follows.

```
# Record 3873, lane 1   1926  # 1
# Normal  12:27:12.114  0.00  1567    1          33.7    4.46       Car
3553.114     48
3553.127     136
3553.140     358
3553.153     580
3553.166     904
3553.179     1228
3553.192     1397
3553.205     1567
```

3553.218 1519

3553.231 1343

3553.244 1167

3553.257 830

3553.270 492

3553.283 336

3553.296 74

3553.309 17

The meanings of each part of the SIG file will be listed first. The number signs on the first two lines indicate the message header portion of each record. The remaining lines comprise the data points which make up the vehicle signature. On the first line, 3873 is a sequential record number assigned by the technician's data collection software. This sequence is in order of the time the vehicle was last detected by the loops, not the order that the vehicles first crossed the loops. Lane 1 is the lane number that the vehicle was travelling in. The next field value 1926 is the matched vehicle ID from the video database. The number 1 following the number sign '#' is the vehicle class determined from the video ground truth. This number is checked with the result of the vehicle classification algorithms to determine the accuracy of the algorithm. A 13 vehicle class scheme listed in Table 8 was used for this number.

Table 8. Vehicle Class Code

Type	**Code**
Passenger Car	1
Minivan	2
Sports Car	3
Sports Utility Vehicle	4
Van	5
Limousine	6
Trucks, 2 axles	7
Vehicle Towing Trailer	8
Bus	9
Full-size Truck	10
Pickup	11
Station Wagon	12
Trucks, > 2 axles	13

The second line has the first field, Normal, which lists whether the vehicle is in one or two lanes. If the vehicle was straddling lanes, the word "DOUBLE" appears instead of "Normal". The next field, 12:27:12.114, is the time stamp of the vehicle when it was first detected. The next field is the time offset 0.00. Following the time offset is the max amplitude value of 1567 recorded from the loop detectors. The lane number 1 is repeated in the next field. Next field is the speed, 33.7, in meters/second for the front of the vehicle. Next is the value 4.46, which is the length in meters. "Car" is the value for the field that is derived from a coarse classification using vehicle length. The classification uses the following lengths:

MotorCycle <3.0m (generally)

Car 3.0m-4.7m

SmallVan 4.7m-6.0m

SmallTruck 6.0m-8.0m

RigidTruck >8.0m

For signatures with two or more peaks:

RigidTrucks <11.0m

Artic >11.0m

The next two fields with the values 208 and 208 are proprietary fields.

The remaining lines on the signature files are the time in seconds and the magnitude. The time is referenced from car number 1 which receives time 0.XXX where XXX represents the decimal fraction of a second in real time. So if car 1 crossed the detector at 11:22:31.545, then the first time would read 0.545, and all other times would be referenced back to 11:22:31.000. Car # 3873 is therefore 3553.114 seconds later than the zero time. Magnitude was recorded every .012 to .014 seconds in most cases. For those vehicles that spent more than 10 seconds over the loop, the time was recorded backwards, recording a number every 0.001 seconds.

For most vehicles a graph of magnitude vs. time can be created. As long as the vehicle was moving at a constant speed or at a gradual change in speed, signatures that represent the vehicle signature can be created from the signature files. If the vehicle stopped or significantly changed speeds over the loops, then irregular signatures may appear.

3.2 DATA PROCESSING

Once the inductive signatures are collected from vehicles, they can be processed to produce useful feature vectors. Feature vectors that were from the inductive signatures are transformed signature, vehicle electronic length, inductive magnitude, energy content, signature shape variance, signature shape skewness, signature shape kurtosis, Karhunen-

Loeve transform coefficients, Discrete Fourier Transform, and the first and second order derivatives of the signature. These feature vectors are not completely independent from each other, so that a careful choice of the vectors can lead to reduction in input data.

There are three main steps in deriving the transformed signature from the raw signature. These steps are also illustrated in Figure 11.

Step1. The signal magnitude is normalized with respect to its maximum amplitude to eliminate variations between detector stations.

Step 2. In order to remove the effects of vehicle speed on the signature, we transform the time axis into length by multiplying by the speed of the vehicle that is obtained from each double loop speed trap configuration.

Step 3. The signature is interpolated so that consistent sample points can be obtained for comparison between different signatures. A spline interpolation (Shoenberg, 1946) was used to fit the signature and provide sampled data points. These sampled ordinates become the values in the signature feature vector. The spline interpolation is a piecewise polynomial interpolation that requires the endpoints of the approximating polynomial to equal the endpoints of the piecewise function. In other words, given a function $f(x)$ on an interval $a\ f(x)\ b$, partitioned into subintervals $a=x_0<x_1<...<x_n=b$, then $g(x_0)=f(x_0)$, $g(x_1)=f(x_1)$, ..., $g(x_n)=f(x_n)$ for all approximating polynomials, $g(x)$, where $g(x)$ should be several times differentiable. Spline interpolation is chosen because it generally avoids the problem of numerical instability and is easy to implement. The derivation of electronic length, inductive magnitude, and energy content feature vectors is straightforward once the raw signature is transformed.

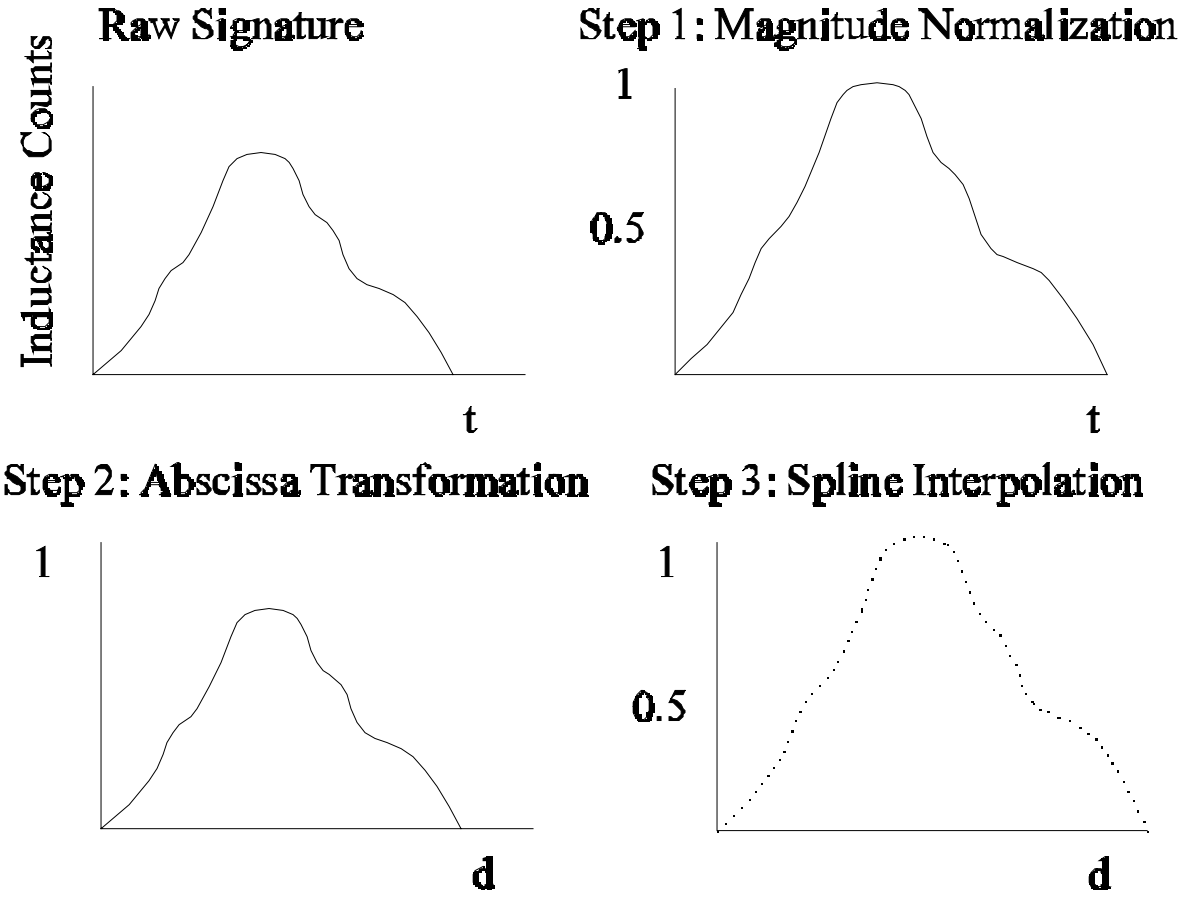

Figure 11. Signature Transformation

The term signature shape is used to emphasize the fact that plotted signatures possess distinguishable shapes. The signature shape variance, signature shape skewness, and signature shape kurtosis are respectively the second, third, and fourth moments of the signature shape. For any random variable X or any function g(X), moments can be defined. Moments are parameters of the variable or function that describe the property of the distribution or function (Kreizig, 1993). Even though moments are often used in the

study of probability distributions, it is equality applicable when analyzing functions or "shapes" such as vehicle signatures.

The first four moments describe the central location, dispersion, asymmetry, and "fatness of the tails". The first two moments, mean and variance are common and are often used when defining normal distributions. For discrete variables, the third and fourth moments are expressed mathematically as

$$\mu_3 = E([X-\mu)]^3 = \sum_j (x_j - \mu)^3 f(x_j)$$

$$\mu_4 = E([X-\mu)]^4 = \sum_j (x_j - \mu)^4 f(x_j)$$

The standardized values of the third and fourth moments are defined as

$$\text{standardized third moment} = \frac{\mu_3}{(\mu_2)^{3/2}}$$

$$\text{standardized fourth moment} = \frac{\mu_3}{\mu_2{}^2}$$

If the signature has a longer tail extending to the "left side" then it is negatively skewed and the standardized third moment is less than zero. If the signature has a longer tail extending to the "right" side then it is positively skewed and the standardized third moment is greater than zero. If the signature has thinner tails than a normal distribution, then it has positive kurtosis and the standardized fourth moment is greater than three. If the signature has fatter tails than a normal distribution, then it has negative kurtosis and the standardized fourth moment is less than three (Bollen, 1989).

The well-known Karhunen-Loeve (KL) method, also known as principle component analysis, has been used in many applications such as speech data compression (Chen and Huo, 1991), neural network training (Malki and Moghaddamjoo, 1991), and image analysis (Abbas and Fahmy, 1993). Some researchers working on vehicle reidentification have used the Karhunen-Loeve expansion to reduce the dimensionality of the vehicle signature vector (Böhnke and Pfannerstill, 1986). In addition to compression, the principle components can also be used as feature vectors of the vehicle signature.

The KL method is based on the creation of a correlation matrix from either time or frequency domain data. The principle components are formed from the eigenvectors and eigenvalues of the correlation matrix. The amount of information that is contained in the principal components is proportional to the magnitude of the associated eigenvalue. Therefore the size of the eigenvalue is an indication of the percent of variance explained by the principle component. Geometrically, the principle components represent the selection of a new coordinate system with the axes representing the directions with maximum variability. In other words, the principal components are the uncorrelated linear combinations of a vector whose variances are as large as possible (Johnson and Wichern, 1992).

Let ρ be the autocorrelation matrix associated with the vehicle signature vector $X'=[X_1, X_2, \ldots, X_N]$. The autocorrelation matrix can be derived using the approximate formula (Chatfield, 1996)

$$\rho_k = \frac{\sum_{t=1}^{N-k}(x_t - \bar{x})(x_{t+k} - \bar{x})}{\sum_{t=1}^{N}(x_t - \bar{x})^2}$$

The variable k is the number of lag intervals. This formula is approximate because it does not consider the degrees of freedom and assumes that N is large. The i^{th} principle component is formed by using the i^{th} eigenvalue/eigenvector pair (λ_i,e_i) of the correlation matrix by $Y_i=e_i'X$, where i=1,2,…N. If a small number of eigenvalues contain much of the variability, then the original system can be represented by the small set of associated principle components without significant loss of information.

Fourier analysis is a well-known technique for performing frequency domain and spectral analysis. Linear time-invariant (LTI) systems are systems that posses the property of superposition and produce a time shifted output when the input is time shifted. Fourier analysis utilizes the properties of LTI systems since the response of an LTI system to any input consisting of a linear combination of basic signals is the linear combination of the individual responses to each of these basic signals (Oppenheim and Willsky, 1983). By using combinations of harmonically related complex exponentials, a signal can be decomposed and the corresponding spectral coefficients can be determined. The spectral coefficient is a term used in spectroscopy in which the spectra is decomposed into spectral lines or frequencies. The intensity at any line is the fraction of the total light energy at the frequency corresponding to the line.

For finite duration signals, the discrete Fourier transform (DFT) is ideally suited for performing frequency domain analysis using digital computers. Given a signal of finite duration with length N, *x[n]*, the DFT synthesis/analysis pair can be expressed as:

$$x[n] = \sum_{k=0}^{N-1} X(k) e^{jk(2\pi/N)n},\ n = 0,1,\ldots,N-1$$

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x[n] e^{-jk(2\pi/N)n},\ k = 0,1,\ldots,N-1$$

where X(k) are the DFT coefficients. If *k=0*, then the D.C. or constant component of *x[n]* can be obtained. This so-called "fundamental" component is the average value of *x[n]*. Since this average value is already captured by other feature vectors, this first component is not used in our graphical analysis of the vehicle waveform spectrum. Because the Fourier coefficients are complex exponentials, graphical techniques are used to simplify the frequency domain analysis. Appendix A contain plots of the vehicle waveform DFT in the complex plane for some representative vehicles. Arrows initiating from the origin show the magnitudes and arguments of the DFT coefficients. For a complex coefficient *X(k)=a+bi*,

$$\left|X(k)\right| = \sqrt{a^2 + b^2}$$

$$\theta = \tan^{-1}\frac{b}{a}$$

Figures in Appendix A show the symmetry property of the Fourier Transform. Since *x[n]* is a real-valued sequence *X(Ω)=X*(-Ω)*, *Re{X(Ω)}* is an even function of *Ω*, and

Im{X(Ω)} is an odd function of *Ω*. Likewise, the magnitude of *X(Ω)* is an even function and the angle is an odd function (Oppenheim and Willsky, 1983). Using this graphical technique, it is clear that the DFT from more unique vehicles such as trucks and buses have higher frequency components than less unique vehicles such as cars and sports utility vehicles.

There is a direct relationship between the Karhunen-Loeve transform and the Fourier transform. The autocorrelation function and the spectral density function contain the same information about intertemporal dependence (Fishman, 1969), and they form a Fourier transform pair. There are trade-offs in performing analysis in the frequency versus the time domain. The autocorrelation function emphasizes the time axis dependencies, while the spectral density functions emphasizes the oscillations at different frequencies. In other words, the spectral density function can be used to describe the relative importance of each uncorrelated frequency component, while the autocorrelation function expresses the weighted sum of the same components. The superposition property makes the DFT additive and the existence of uncorrelated coefficients lends itself to graphical methods.

The first and the second derivatives of the signature are useful in determining the number of maxima and minima. The usefulness of the extreme points is most easily understood through the visual inspection of the signature plots. Signatures of passenger cars are almost always parabolic and contain only one maximum and no minima. More unique

vehicles contain multiple maxima in their signatures. See appendix A for examples of signature plots and the corresponding feature vectors.

CHAPTER 4. DECISION THEORETIC VEHICLE CLASSIFICATION

The process of classifying vehicles can be divided into two phases: feature extraction and classification. The feature extraction portion seeks to extract the salient components of inductive signatures that would sufficiently differentiate vehicle classes. In order to avoid redundancy, each vector would contain different information. This is similar to the process of deriving a basis in linear algebra. In a likewise fashion, the goal here is to find an orthogonal set of vectors that would span the pattern space of possible vehicles signatures. Since this classification system is developed with future real-time implementation in mind, issues such as the cost of extraction, storage requirements, and communications bandwidth are of importance in addition to the classification rate.

4.1 HEURISTICS

The heuristic algorithms presented in this paper use select feature vectors that were obtained from processing the inductive signatures. The classification phase involves the use of a heuristic, or a combination of several discriminant functions processed sequentially.

One advantage of this heuristic approach is the reduction of data and communication requirements by using processed feature vectors. A fixed number of feature vectors will be inputs to the classification algorithm instead of entire signatures. As described in the previous section, the feature vectors that were considered include the vehicle electronic length, inductive magnitude, energy content, signature variance, signature skewness,

signature kurtosis, Karhunen-Loeve transform coefficients of the signature, Discrete Fourier Transform of the signature, and number of maxima/minima of the signature. Only a subset of these feature vectors were used in developing the three heuristic algorithms as some feature vectors contained redundant information. Other feature vectors that were not used in the three heuristic algorithms presented here can possibly be used efficiently in other classification algorithms related to particular transportation applications.

Another advantage of the heuristic approach is that the sequential classification stages are explicitly defined. Therefore problems with a particular stage can be identified and corrected without having to rebuild a complete classification algorithm. The sequential aspect of the algorithm also results in efficient computation since the feature space is reduced as the algorithm proceeds sequentially.

The following paragraphs discuss the classification portion of the heuristic algorithm. One approach in pattern recognition algorithms is to use a nonparametric decision theoretic approach. In this approach, no assumptions of class distributional forms are made. If each signature is considered to be a pattern in the pattern space, then signatures pertaining to different vehicle classes will fall into different regions of this pattern space. A separating surface can then be defined to differentiate between the different classes. In one-dimensional space, the decision surface is a point that divides the linear pattern space as seen in Figure 12 (Bow, 1992).

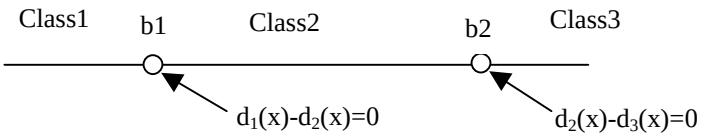

Figure 12. One-dimensional Pattern Recognition Example

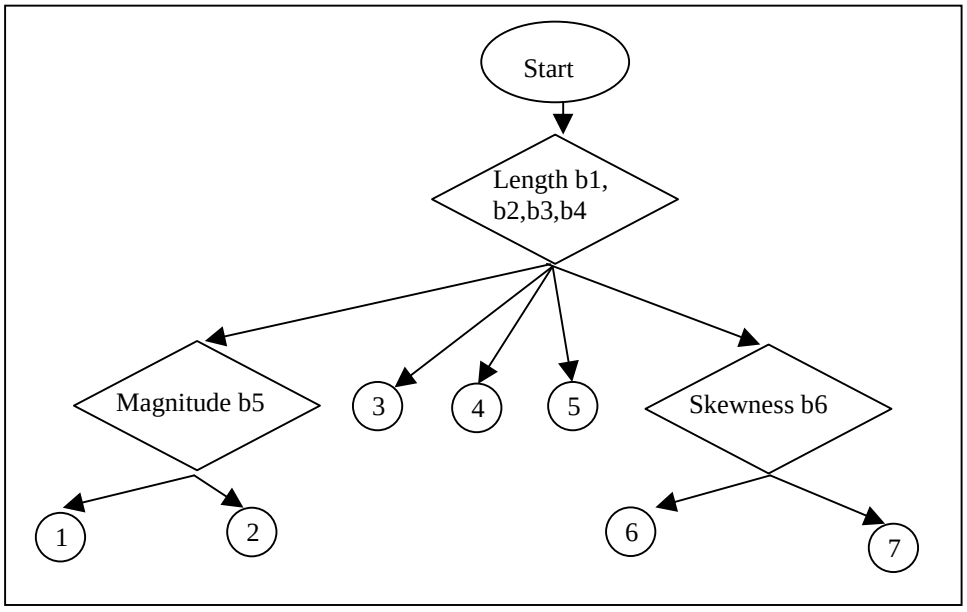

Figure 13 (a). Heuristic Algorithm 1

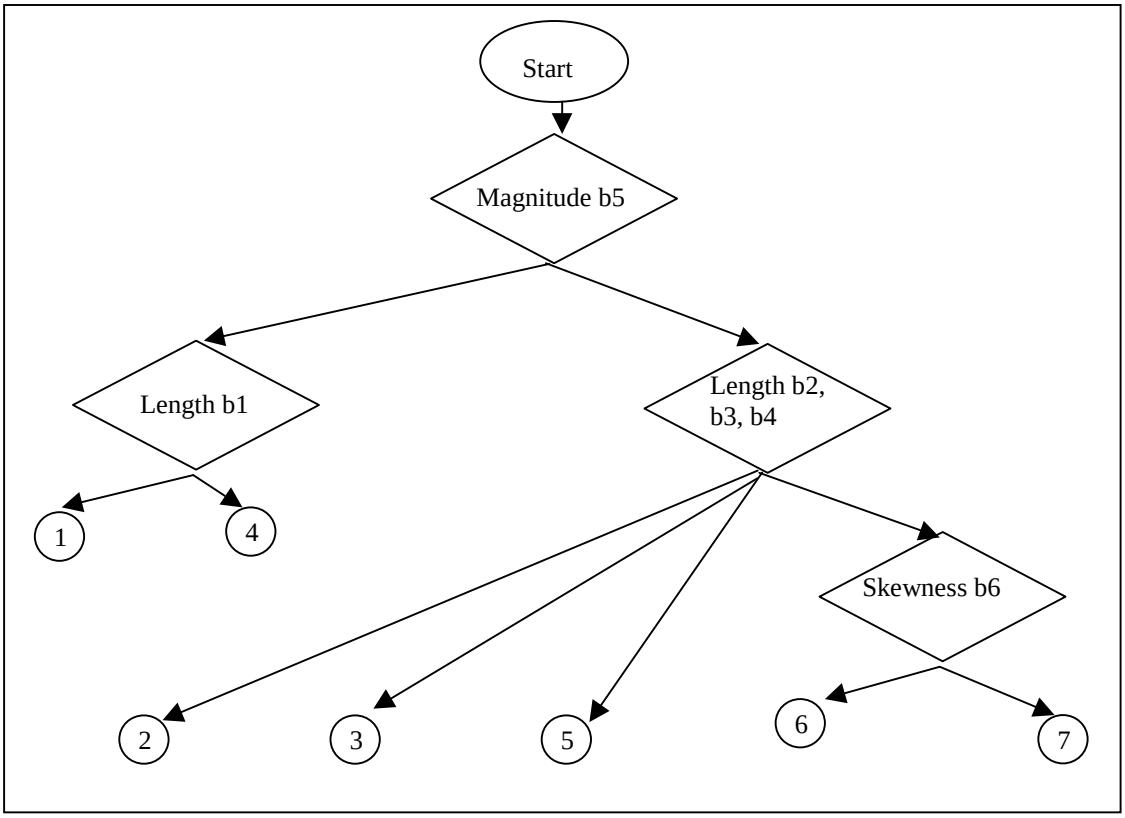

Figure 13 (b). Heuristic Algorithm 2

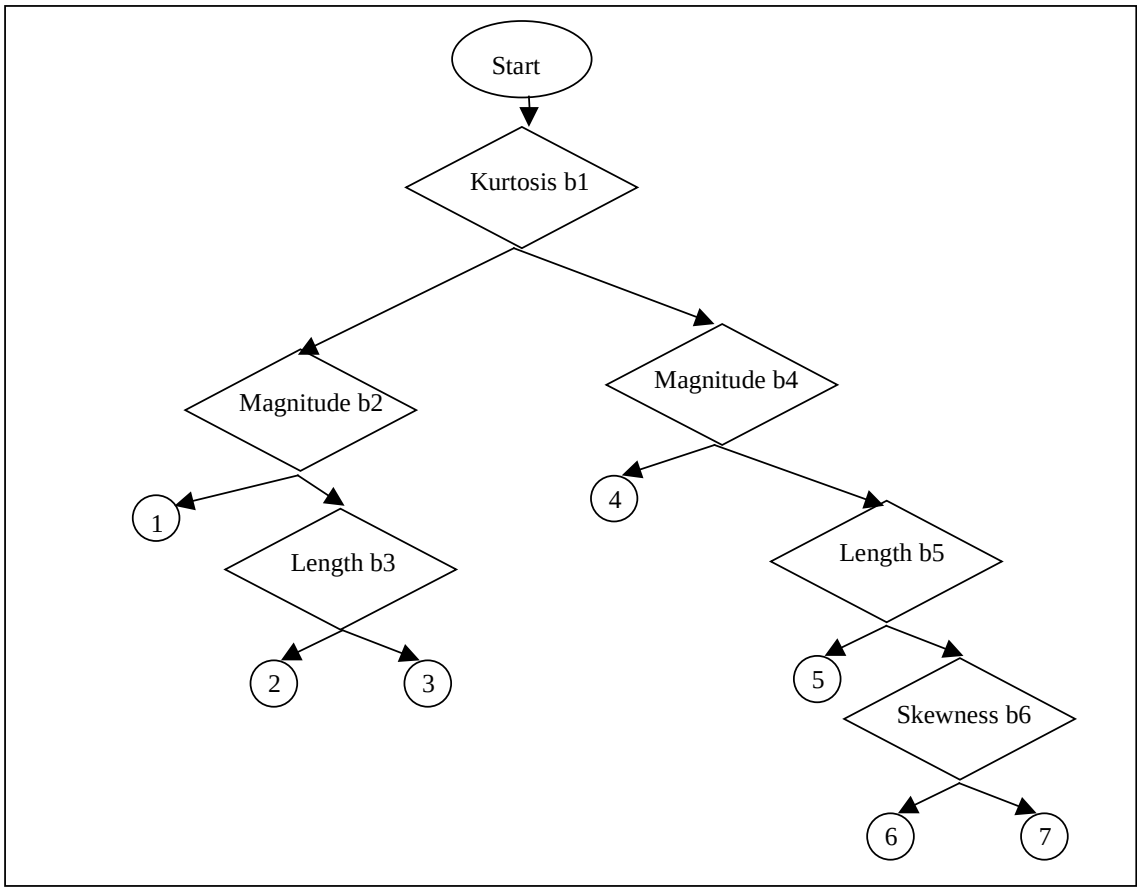

Figure 13 (c). Heuristic Algorithm 3

Assume that the use of each feature vector is a one-dimensional discrimination problem, then vehicle classification can be performed by using a single feature vector. However, since the feature vectors are not completely correlated or redundant, multiple feature vectors can be combined to improve classification rate and to increase the number of classification categories. Conceptually, this is illustrated in Figure 13 in which heuristic classification trees use several feature vectors to arrive at seven separate classes. These heuristic classification algorithms which combine multiple feature vectors is the approach undertaken for this research. Figure 14 shows an hierarchical structure that is similar to

the hierarchical structure of syntactic pattern decomposition schemes and is equivalent to a piecewise linear discriminant function.

4.2 TRAINING METHODOLOGY

The training of the heuristic algorithms involves the determination of the optimum discriminant bounds of the discriminant functions used in the heuristic tree. The objective in training the heuristic algorithm is to minimize the number of misclassifications or

$$\min_{\tilde{b}} \sum_{j=1}^{N_v} \delta(\tilde{b}, \tilde{x}_j), \quad \forall i$$

where $\delta(\tilde{b}, \tilde{x}j) = \begin{cases} 1, \text{ if } x_j \text{ is classified correctly} \\ 0, \text{ if } x_j \text{ is classified incorrectly} \end{cases}$, N_v is the number of training vectors, $\tilde{b}$ is a vector of boundaries that define the decision surfaces. Note that each training vector derived from a single vehicle signature is composed of multiple feature vectors. In other words,

$$\tilde{x}_j = x_{j1}, x_{j2}, ..., x_{jN_f}$$

and N_f is the number of feature vectors per vehicle signature.

This optimization problem can be reformulated as a hierarchical multi-objective optimization (or lexicographic optimization) problem. Instead of finding the components of $\tilde{b}$ simultaneously, a staged approach is taken to find each component of

$\tilde{b}$ sequentially. At each stage, a separate feature vector is considered, and the decision surfaces are determined. A well-known minimization technique called golden section search was used for each optimization stage. The optimization is complete when Pareto optimality is achieve. At such an instance, it is not possible to move feasibly so as to increase one of the objectives without decreasing at least one of the others (Steuer, 1986).

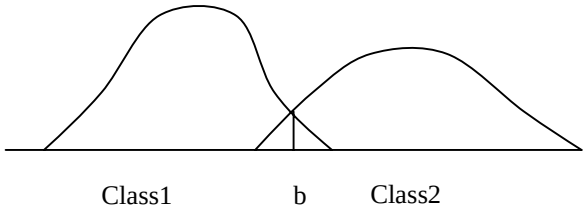

Figure 14. Example of Discrimination with Unimodal Feature Vector Distributions

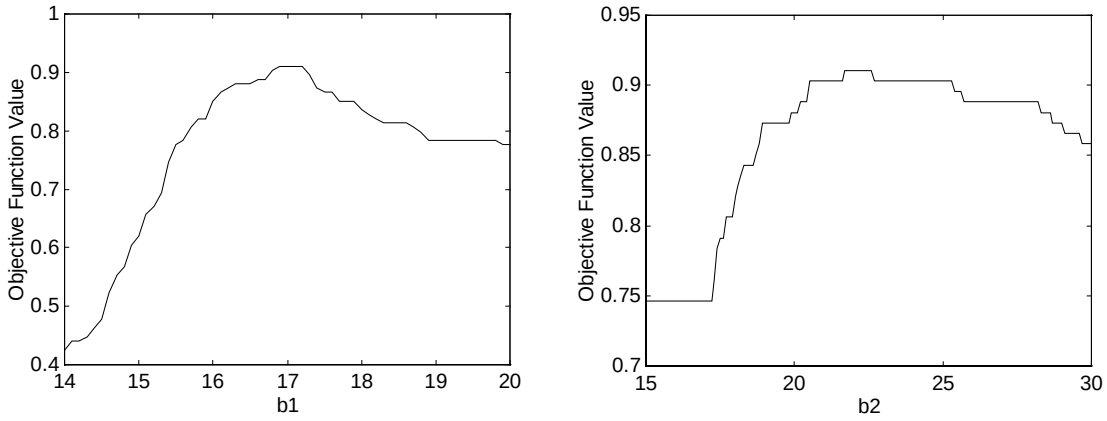

Figure 15. Examples of Unimodal Objective Functions Using Vehicle Signature Length as Feature Vector

Golden section search was used because the only critical property for applying this method was for the objective function to be unimodal (Luenberger, 1989). Since the objective function is derived heuristically it has no mathematical functional form, thus

gradient search methods can not be used. The unimodal property is satisfied when the distributions of the feature vectors are unimodal. Figure 14 illustrates a simple example of unimodal feature vector distributions for two classes. This property needs to be checked before a feature vector is used. If a distribution is multi-modal, then an inferior solution can result from searching and finding a local minimum. Figure 15 gives examples of the unimodal objective function from using the length feature vector which has a unimodal distribution.

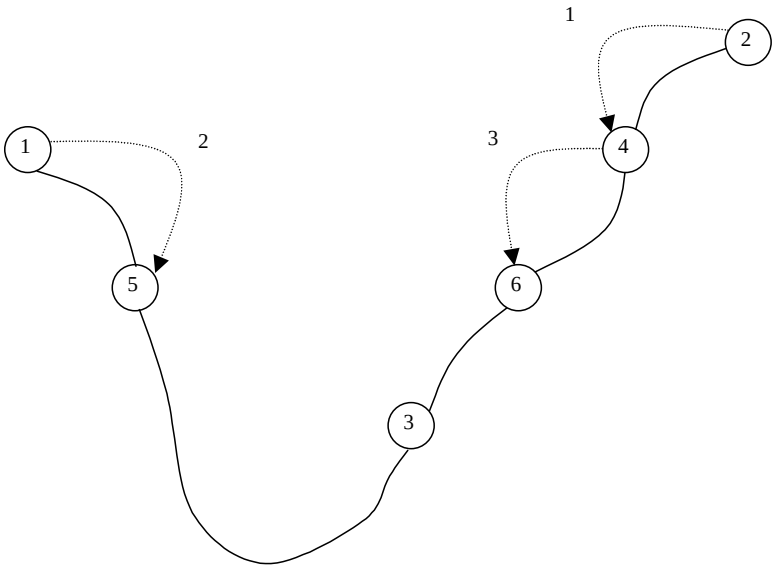

Figure 16. Graphical Illustration of Golden Section Successive Bracketing

Golden section search is a method which brackets the minimum of an objective function at a certain tolerance level. A minimum is bracketed by three points (a,b,c) such that *f(b) < f(a), f(c)*. Initially, the bracket is set arbitrarily, and the bracket is incrementally reduced as new evaluations of the objective function become available. Figure 16 illustrates this incremental bracketing of the minimum. In Figure 16, the initial bracket is [1,3,2]. Then the objective function is evaluated at point 4 and a new bracket becomes [1,3,4]. This procedure continues by always keeping the center point lower than the outside points (Numerical Recipes, 1992). Golden Section Search derives its name from using the optimum bracketing distances of 0.38197 and 0.61803 known as the golden section. This golden section is a result of scale similarity between each successive bracket.

In practice, the initial bracketing points in this algorithm is set by first examining the feature vectors of 26 representative signatures from different classes. If the initial bracket does not contain the minimum, then the search procedure will end up finding the bracketing point as the minimum, at which time the bracket should be enlarged. This linearly converging search method results in computation times of only a few minutes for training the classification algorithms.

4.3 RESULTS

Three different heuristic algorithms were developed and tested. Conceptually, the three heuristics are illustrated in flowcharts in Figure 13. These three algorithms differ in the feature vectors that are used, in the order in which they are used, and in the number of

discriminant surfaces per feature space. As seen in Figure 13, the first level discriminants used for the three algorithms were length, magnitude, and kurtosis. Algorithms one and two are more similar because they use multiple discriminant surfaces for the length feature space, while the third algorithm only makes binary decisions.

The purpose in presenting three algorithms is to show examples of how different feature vectors can be combined in various ways to produce classification algorithms. This research was not intended to promote a single algorithm nor to promote an optimum set of feature vectors for all applications. The three algorithms then are used to demonstrate the feasibility of using the heuristics for vehicle classification. For a particular ATMIS application, the heuristic algorithms can be modified to suit that application.

Table 9. Seven Vehicle Class Scheme

Class Code	1	2	3	4	5	6	7
Description	car, minivan, sports, station wagon	SUV, pickup	van, full-size pickup	limo	2 axle truck	vehicle + trailer, bus	> 2 axle truck

A vehicle classification scheme using seven vehicle classes was used for testing algorithms. This scheme is described in detail in Table 9. This scheme is chosen because it yields classes that can be adapted for several applications including pollution, safety, and traffic analysis. This scheme is also chosen because it targets vehicle classes that are

not differentiable by using current techniques based on axle counting. For any particular application, the number of classes in the heuristic classification scheme can be decreased or increased to suit that application.

A total number of three hundred vehicle signatures, including ground truth video, were collected for this investigation. The manageable size of the data set allows detailed analysis of each vehicle signature and the examination of the causes for each individual misclassification. The field data were obtained from two freeway sites on the westbound SR-24 freeway in the city of Lafayatte in Northern California. By training and testing using data collected from different sites, the transferability of the algorithms was tested. Note that no calibration of the loops was performed between sites or within the multiple loops at each site. The distance between the two freeway sites was approximately 2 km (1.2 miles) long. At each detector station there were four lanes. Each data acquisition station was instrumented with video, inductive signature dataloggers, and speed trap dataloggers. Standard 1.82mX1.82m (6'X6') loops were used at both stations. The data set was separated into a training set and a test set of roughly equal size. The number of vehicles in each of the classes was not equal, thus there were more vehicles in the first three classes than in the last four. The reduced data set contained the signatures of the vehicles along with their speeds, arrival time at stations, lane number, and the ground truth video. This dataset was composed of moderate flow traffic (~1000 VPHPL). The calibration of the algorithms was performed on the training data only.

Table 10. Optimized Discriminant Bounds for Heuristic Algorithms

Algorithm	b1	b2	b3	b4	b5	b6
Heuristic 1	15.94	19.97	26.37	33.01	1000	10
Heuristic 2	16.87	16.77	21.68	29.67	1019.13	2.24
Heuristic 3	21.729	983.172	15.910	1250.1	29.438	10

Table 10 shows the heuristic bounds used in each algorithm. These bounds were determined by using golden section search. The tolerance used in stopping the golden section search was for the distance between the outside bracket points to be smaller than 0.0005.

Table 11. Classification Results

Algorithm/ Data	Overall Rate	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7
Heuristic 1/Training	84%	86%	92%	81%	100%	50%	100%	75%
Heuristic 1/ Testing	81%	87%	83%	73%	100%	50%	100	67%
Heuristic 2/ Training	91%	96%	100%	77%	100%	75%	100%	50%
Heuristic 2/ Testing	82%	87%	96%	69%	67%	75%	50%	50%
Heuristic 3/ Training	88%	93%	88%	77%	67%	88%	100%	75%
Heuristic 3/ Testing	85%	96%	72%	77%	67%	75%	100%	75%

The overall classification rates from using heuristic 1 were very similar between the training and the test set as seen in Table 11. Algorithm 1 is heavily dependent upon the length feature vector, since the length is used as the top level discriminant and it is used to differentiate five categories. Even though length proved to be effective as a top level discriminant, its use was not sufficient to separate class 1 from class 2 or class 6 from

class 7. The inductance magnitude was used to separate minivans and SUV's of the same length, and skewness was used to separate trucks and buses of the same length. One advantage of using inductive signatures over length as an input to a classification system is the ability to separate vehicle classes among vehicles of the same length. The performance of algorithm 1 was good for all classes except class 5. This motivated the development of another algorithm that would maintain the overall classification rates while at the same time improve performance for class 5.

Heuristic 2 was developed by replacing length with magnitude as the top level discriminant. The basis for this decision was the fact that the inductance magnitude is approximately inversely proportional to the square of the height of the vehicle undercarriage from the loop. Therefore trucks and other vehicles with high undercarriage clearances would produce signatures with lower magnitudes. Even though algorithms 1 and 2 used the same feature vectors, the results were different as seen in Table 11. The classification rate for class 5 increased but classification rates for class 6 and 7 decreased. This was not surprising since class 5 and class 7 are closely related.

For heuristic 3, Table 11 gives the overall classification rate and rates by vehicle class. The overall classification rate of almost 90% demonstrates the potential of using heuristics for vehicle classification. The performance across all categories is more consistent than previous algorithms. The poorest performing category is class 4, composed of limousines. This category is problematic because limousines have inductive magnitudes that are similar to passenger cars while their lengths are similar to trucks.

The results of the test set shows that the classification rate went down slightly as compared to the training set. The individual class rates were also consistent across the board. The results of the test set again point to the potential of using heuristics for vehicle classification.

Table 12. Classification Results for Heuristic 3

		Heuristic 3 Classification for Training Set						
		Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7
Correct Class	Class 1	64	4	1	0	0	0	0
	Class 2	2	22	1	0	0	0	0
	Class 3	4	2	20	0	0	0	0
	Class 4	0	0	0	2	1	0	0
	Class 5	0	0	1	0	7	0	0
	Class 6	0	0	0	0	0	2	0
	Class 7	0	0	0	0	1	0	3
		Heuristic 3 Classification for Test Set						
		Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7
Correct Class	Class 1	66	3	0	0	0	0	0
	Class 2	1	18	6	0	0	0	0
	Class 3	3	2	20	0	1	0	0
	Class 4	0	0	0	2	1	0	0
	Class 5	0	0	1	0	6	1	0
	Class 6	0	0	0	0	0	2	0
	Class 7	0	0	0	0	1	0	3

In Table 12, the results of individual classifications were tabulated. This is useful in analyzing both type I and type II errors in classification. If the null hypothesis is that the vehicle was correctly classified, then Type I error is the rejection of the null hypothesis when it is true and Type II error is not rejecting the null hypothesis when it is false.

CHAPTER 5. SELF-ORGANIZING FEATURE MAP

5.1 NEURAL NETWORKS

A good taxonomy of neural networks for pattern classification is shown in Figure 17. (Lippman, 1987). Classical classification algorithms which are most similar to the neural network is listed in parenthesis. Although there are other taxonomies for neural networks, Figure 17 is helpful in identifying the type or types of neural network(s) needed for a particular application in pattern recognition.

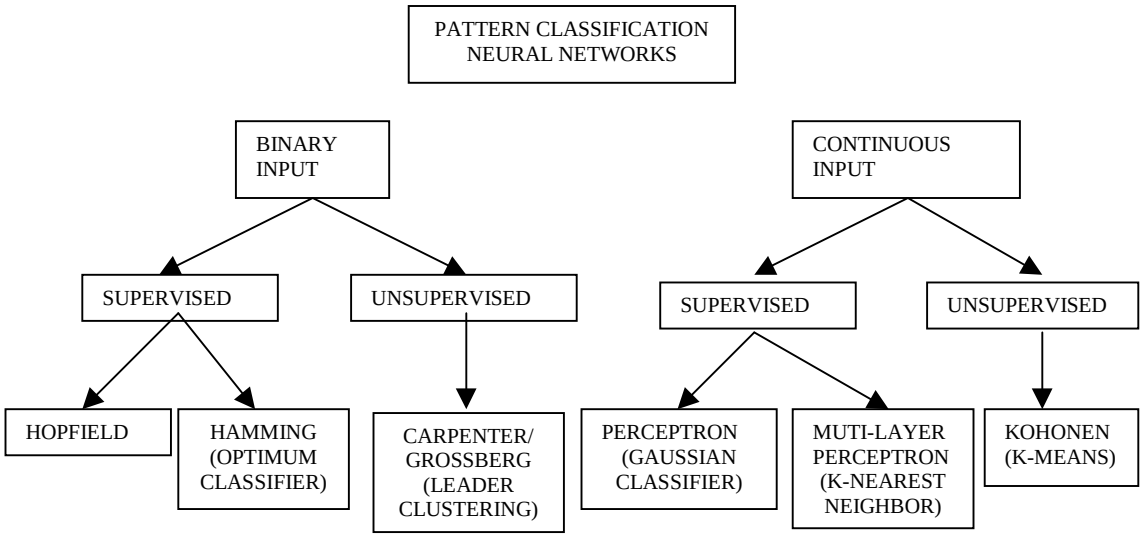

Figure 17. Taxonomy of Pattern Classification Neural Networks

A Self-Organizing Feature Map (SOFM) is an Artificial Neural Network that forms clusters of neurons that reflect similarities in the input vector. SOFM is also known as the Kohonen network. Some references on this subject include (Hecht-Nielsen, 1990) and (Wasserman, 1989). Since SOFM uses unsupervised learning, it is a mapping that is

defined implicitly and not explicitly. Input vectors are presented sequentially in time without specifying the output. Because of this fact, there is no way of predicting which neuron will be associated with a given class of input vectors. This mapping is accomplished after the training by testing the network.

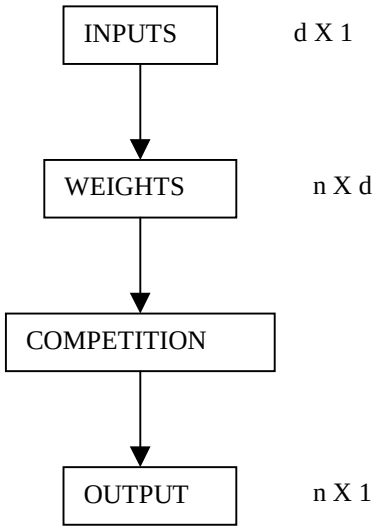

Figure 18. Graphical Representation of the Self-Organizing Feature Map

Figure 18 shows the architecture of SOFM graphically. The d input vectors **x** are received by the n neurons in parallel and are scaled by the weight vector **w**. Thus, the weight matrix is the size of n neurons by d inputs. The n neurons are then entered into competition where only one neuron wins.

SOFM employs the concept of topological neighborhoods, which are equidistant neuron neighborhoods centered around a particular neuron. This concept is better illustrated graphically as in Figure 19 for a two-dimensional case. The actual SOFM's employed in

the classification of vehicle signatures are only one-dimensional. The neighborhood distance matrix for a one-dimensional case in shown in Figure 20 for a network with four neurons. It can be seen in Figure 20 that the distance of a neuron from itself is 0, the distance of a neuron from its immediate neighbor is 1, and so on. Unlike simple competitive learning, the weights of the neighborhood neurons are updated in addition to the weights of the winning neuron.

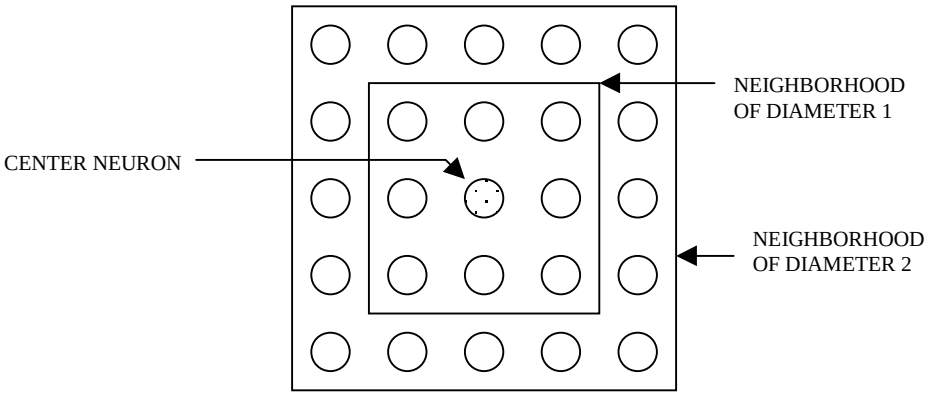

Figure 19. Topological Neighborhoods

0	1	2	3
1	0	1	2
2	1	0	1
3	2	1	0

Figure 20. Topological Neighborhood Distance Matrix

The steps in the training of the SOFM can be outlined as follows:

Step 1. Initialize weights randomly.

Step 2. Present new input.

Step 3. Compute distances between input and neuron weights.

Step 4. Competitive selection of the neuron with the minimum distance.

Step 5. Update neuron weights to winning neuron and neighborhood neurons.

Step 6. Continue iterating by going through steps 2-5.

In describing the training algorithm, it is useful to understand the weight structure of the network. Each neuron has the same number of weights **w** as the dimension of the input vector **x**. The weight structure for each neuron can then be viewed as a matched filter competing against other neurons. A matched filter has the impulse response tuned to input so that it produces the maximal output signal (Oppenheim, 1983). The overall weight structure can be viewed as an array of matched filters with each neuron's weights being adjusted on the basis of current weights and the goodness-of-match of the input.

Given that the feature map is initialized with random weights $\boldsymbol{w}$ for all the neurons, a distance measure, $d(\boldsymbol{x},\boldsymbol{w_i})$ can be used to define the goodness-of-match between a particular weight vector $\boldsymbol{w_i}$ and the input vector $\boldsymbol{x}$. The distance measure used can be correlation, Euclidean, city-block or other statistical measures. If the Euclidean distance is used, then the winning neuron from the competition is found at each iteration k by using the following equation

$$\|x(k) - w_c(k)\| < \|x(k) - w_i(k)\|, \forall i$$

where neuron c is the winning neuron. The network weights are then updated as follows:

$$w_i(k+1) = \begin{cases} w_i(k) + \alpha(k)[x(k) - w_i(k)], i \in N_c(k) \\ w_i(k), i \notin N_c(k) \end{cases}$$

where N_c defines the neuron neighborhood and $\alpha(k)$ is the learning rate at each iteration. The weight updating equation shows that $d(\mathbf{x},\mathbf{w}_i)$ is decreased for N_c while those outside N_c are left unchanged. The neighborhood size and the learning rate both decrease with increase in the number of training iterations. Typically $\alpha(k)$ start near 1 and go down to 0.1. The size of N_c can start out as large as the greatest distance between weight vectors, and decrease until the neighborhood defines only one neuron.

An interesting feature of SOFM is that the distances between neurons can be interpreted as statistical frequency distributions. In a fully trained network, if input vectors occur with varying frequencies, then the feature map will also allocate neurons to an area in proportion to the frequency of input vectors (Demuth and Beale, 1997).

Another interesting feature of SOFM is that the weights become topologically similar to the input vectors. In a sense, the final weights are analogous to class templates and can be plotted for visual inspection. Because of this feature, the SOFM is less like a black-box and is more like an automated k-means statistical clustering algorithm.

Pursula and Pikkarainen (1994) had already used SOFM for vehicle classification using inductive signatures. The inputs of their system are inductive vehicle signatures with a sampling period of 5-20ms. The physical inductive loop dimensions are 3x2m. The signatures are preprocessed to a standard form in which the x-axis range is uniform across all signatures. In other words, the vehicle length information is removed from the signature. The input dimensions are reduced by considering only the portion of the

signature that lies above the 50% level of its maximum amplitude. Only twenty points are used to describe this upper portion of the signature. Individual vehicle speeds are computed using the 50% amplitude point of the double loop signatures. This speed is used to derive vehicle lengths.

The classification algorithm is composed of three components. The first component is a 12x12 2-dimensional self-organizing feature map, although other network sizes were also used. The second component is a Learning Vector Quantization (LVQ) algorithm that strengthens the correct classifications by a supervised learning process. The final component is a vehicle length window that limits weights vectors to be in the given length range. For the training of the SOFM, 1400 vehicles were used. The testing was accomplished with a test set of 737 vehicles.

The choice of the particular SOFM configuration is not a precise science and involves engineering judgement. The two major design parameters are the dimension of the network and the number of neurons. As noted by Schalkof (1992), powerful results have been obtained by just using 1- and 2-dimensional topologies. In selecting the dimensionality of the network for vehicle classification implementation, the computational intensity and the classification rate are both important. Pursula and Pikkarainen (1994) employed 2-dimensional networks for vehicle classification. They presented encouraging results from the use of 5*5 to 12*12 feature maps. However, a 1-dimensional feature map can also be applied to the vehicle classification problem using less computational intensity. As the vehicle signatures are interpolated at fixed intervals,

there is no need to capture the abscissa. Therefore a 1-dimensional feature map can be used to capture the inductance magnitude changes of the vehicle signatures.

The training parameters are the learning rate and the number of training iterations. The initial learning rate should be between 0 and 1, and value of 0.99 was chosen based on experience. Kohonen (Schalkoff, 1992) cites the use of 10,000-100,000 training iterations as typical, and recommends that the number of training cycle should be at least 500 times the number of output neurons. Thus, all the networks have been trained using 100,000 iterations. By looking at the plot of the feature maps for the different networks used for vehicle classification, one could see that the neurons are well-spaced indicating that the training iterations were adequate. Otherwise, the neurons would be bunched up towards the middle of the input vector space where the training originally started.

5.2 RESULTS

All the feature maps tested had the following general characteristics. First, the input vector is composed of ninety-three components corresponding to the ninty-three equally spaced interpolations of the vehicle signature. Second, the number of training iterations were the same for all feature maps. Even though some trials involving 3000 iterations were performed, the resulting feature maps did not exhibit convergent characteristics. Therefore, based on the recommendations of other researchers and on experience, the number of training iterations was chosen to be 100,000. A 300Mhz personal computer was used for implementing the SOFM, and the training took only a few minutes for each feature map configuration. Third, the training set is composed of only thirteen vehicles

signatures. The 13 vehicles chosen represented one from each of the following categories: passenger car, station wagon, minivan, sports car, Sports Utility Vehicle (SUV), cargo van, limousine, pickup, full-size truck, vehicle towing trailer, bus, box truck, and truck and trailer with more than two axles.

The first implementation of the SOFM involved a 1-D feature map composed of four neurons. Therefore a maximum of four vehicle classes could be derived from this feature map. Four distinct clusters resulted from the training. The four resulting classes (clusters) include vehicle types listed in Table 13.

Table 13. Four Class Scheme

Class	Vehicle Types Represented
1	passenger car, minivan, sports car, station wagon
2	SUV, full-size truck, pickup
3	van, limousine
4	cargo truck, vehicle with trailers, bus, and trucks with >2 axles

One advantage of using SOFM is that the weights actually represent the categorized inputs and can be interpreted as a class template. This makes the graphical analysis of the SOFM possible, and gives insight to the so-called neural network blackbox. The weights of the four neuron SOFM are plotted in Figure 21.

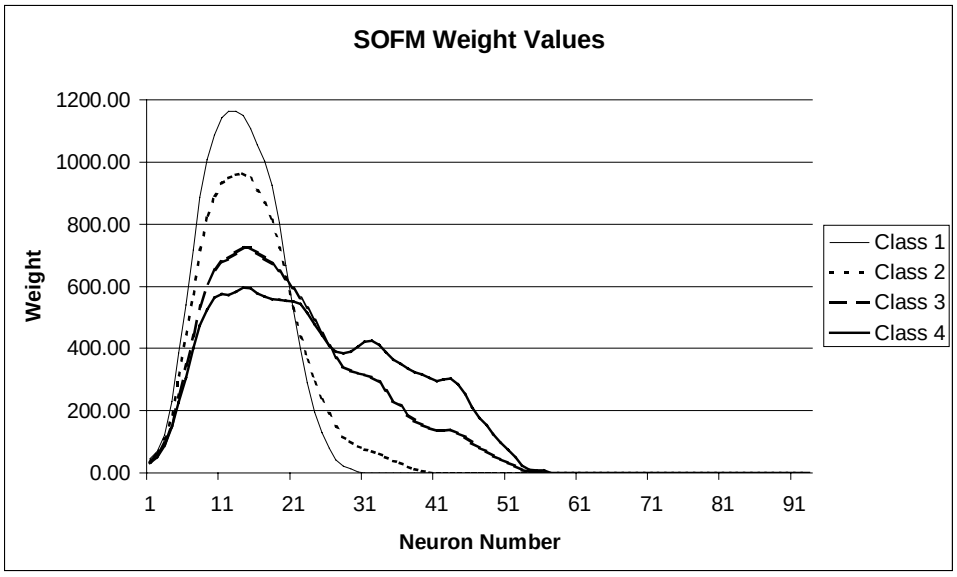

Figure 21. Plot of SOFM Weights

Figure 21 shows that inputs ~55 through 93 are unused. This is a direct consequence of using vehicle length as the x-axis since most vehicles do not require as many input components as compared with a trailer-truck. On the other hand, if the x-axis is normalized, then all input weights would have non-zero values, but the length information would not be captured directly with this SOFM. In such a scheme, the length feature can be incorporated after the SOFM classification with a discriminant function, or with another classification neural network such as a MLF.

A test set composed of 124 vehicles was used. The overall classification rate was 77%. This result is auspicious in view of the small size of the SOFM and the training set. All of the missed classifications incorrectly assigned to an adjacent category. This is encouraging since this implies that a shifting of the classification boundaries might improve the classification rate. The shifting of the classification boundaries is

accomplished by re-training the SOFM. This is perhaps one disadvantage of using neural networks, since these boundaries are not specified explicitly as in statistical approaches but are embedded in the neuron weights. A more detailed analysis of the results show that ~5% of the misclassifications were due to full-size trucks incorrectly classified as class 3.

Detailed analysis of the signatures from the misclassified vehicles was performed. The first group of misclassifications involves cargo vans incorrectly classified as class 2. By analyzing the signature of the training set for the vehicle type “cargo van”, it was found that the van was an ambulance which had a signature that had features of both a van and a truck. Therefore, the training set was modified by replacing the ambulance with a cargo van that was more representative of the signature of the cargo van. Figure 22 shows examples of vehicle signatures of the misclassified cargo vans.

id31 veh47 ar3000 am3.75 an0.58 ev0.47 var1.84 sv0.26 sc0.10 mph73.80 ft16.53 mag799 l2

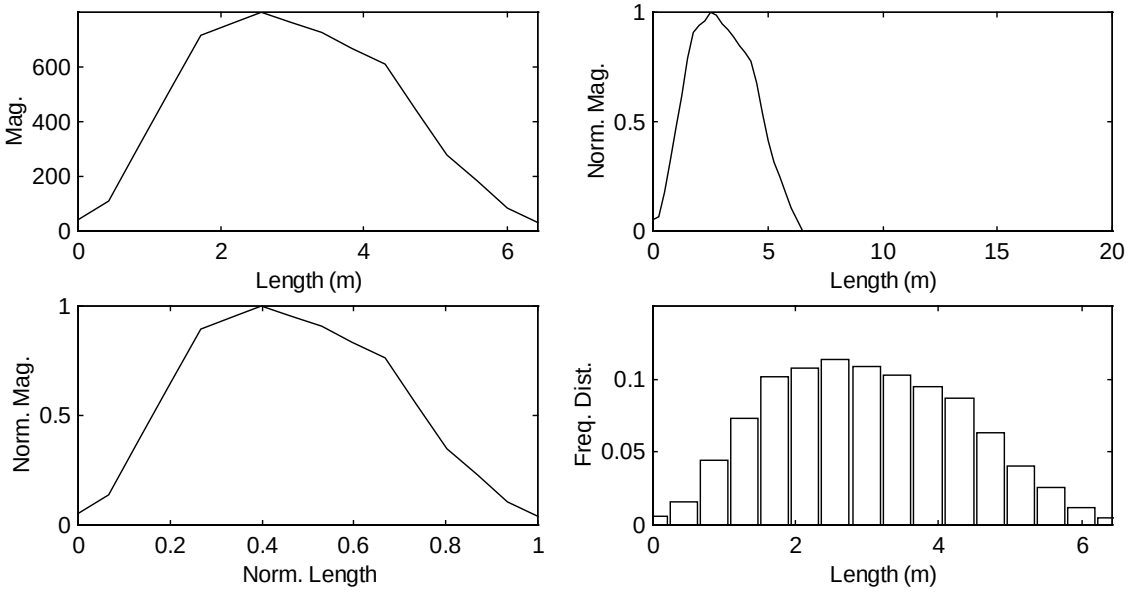

id91 veh140 ar2978 am3.77 an0.51 ev0.48 var1.90 sv0.30 sc0.11 mph70.88 ft15.84 mag790 l2

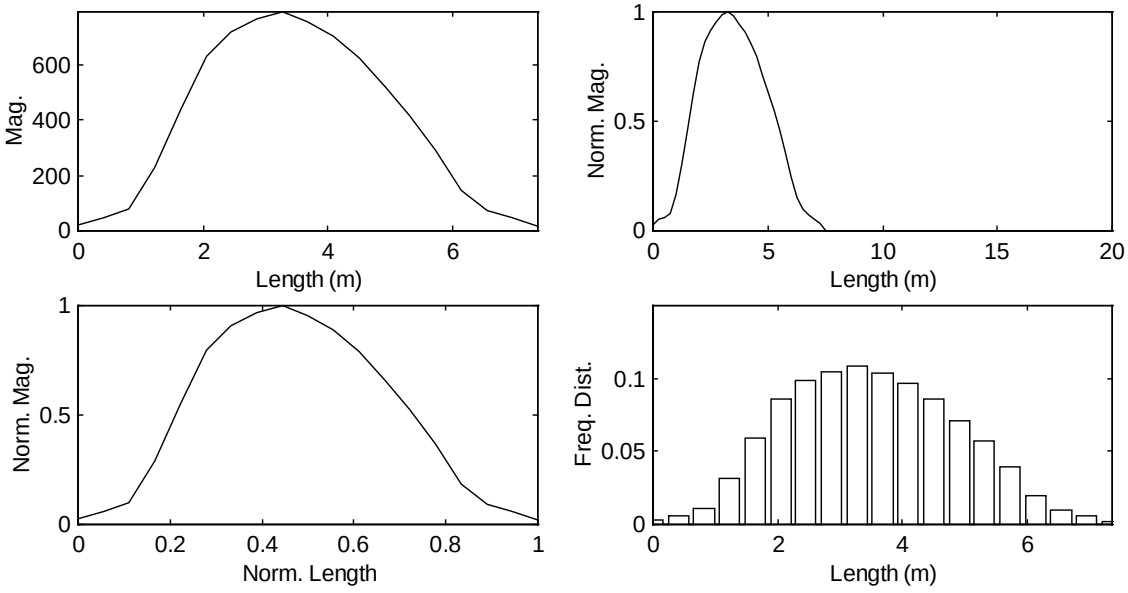

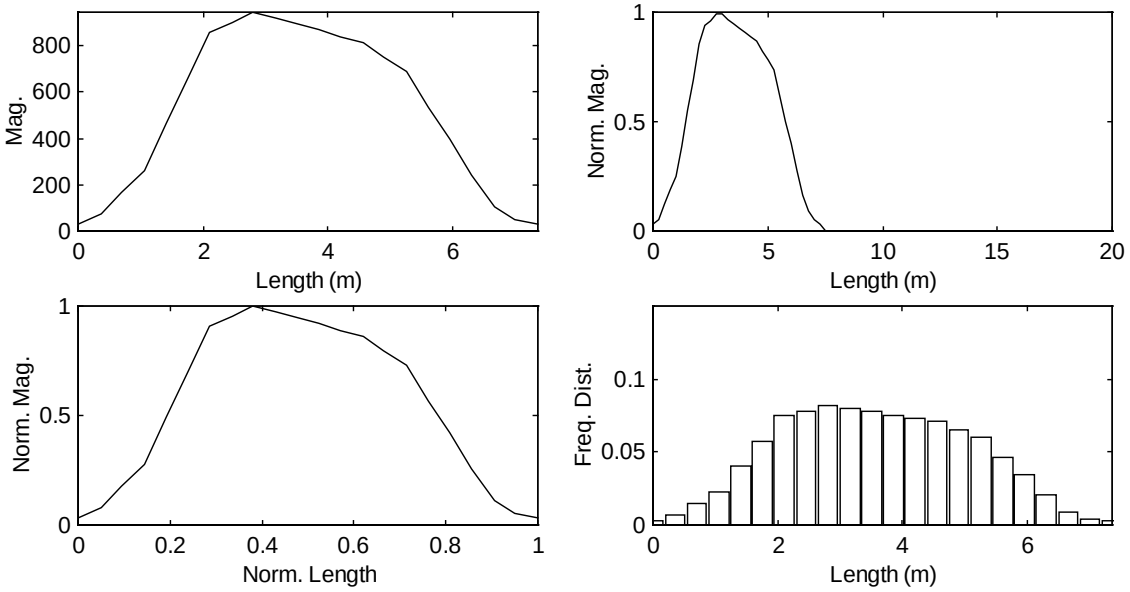

Figure 22. Examples of Misclassified Cargo Vans

Another group of misclassifications involves passenger cars being incorrectly classified into group 2. Figures 23 shows examples of the signatures that fall into this category. Vehicles 98 and 168 are a passenger car and a sports car. Looking at the signatures, one sees that the magnitude of the signature is very low compared with a normally centered passenger vehicle. This indicates that the vehicles were traveling off the center of the lane or switching lanes. Since the inductance change is proportional to the inverse of the distance from the loop detector, pickups and SUV's also exhibit such low magnitude signatures. These two examples demonstrate one weakness of using the signatures from the current inductive loop structure. Namely that lateral offset causes low magnitudes in vehicle signatures which leads to misclassifications.

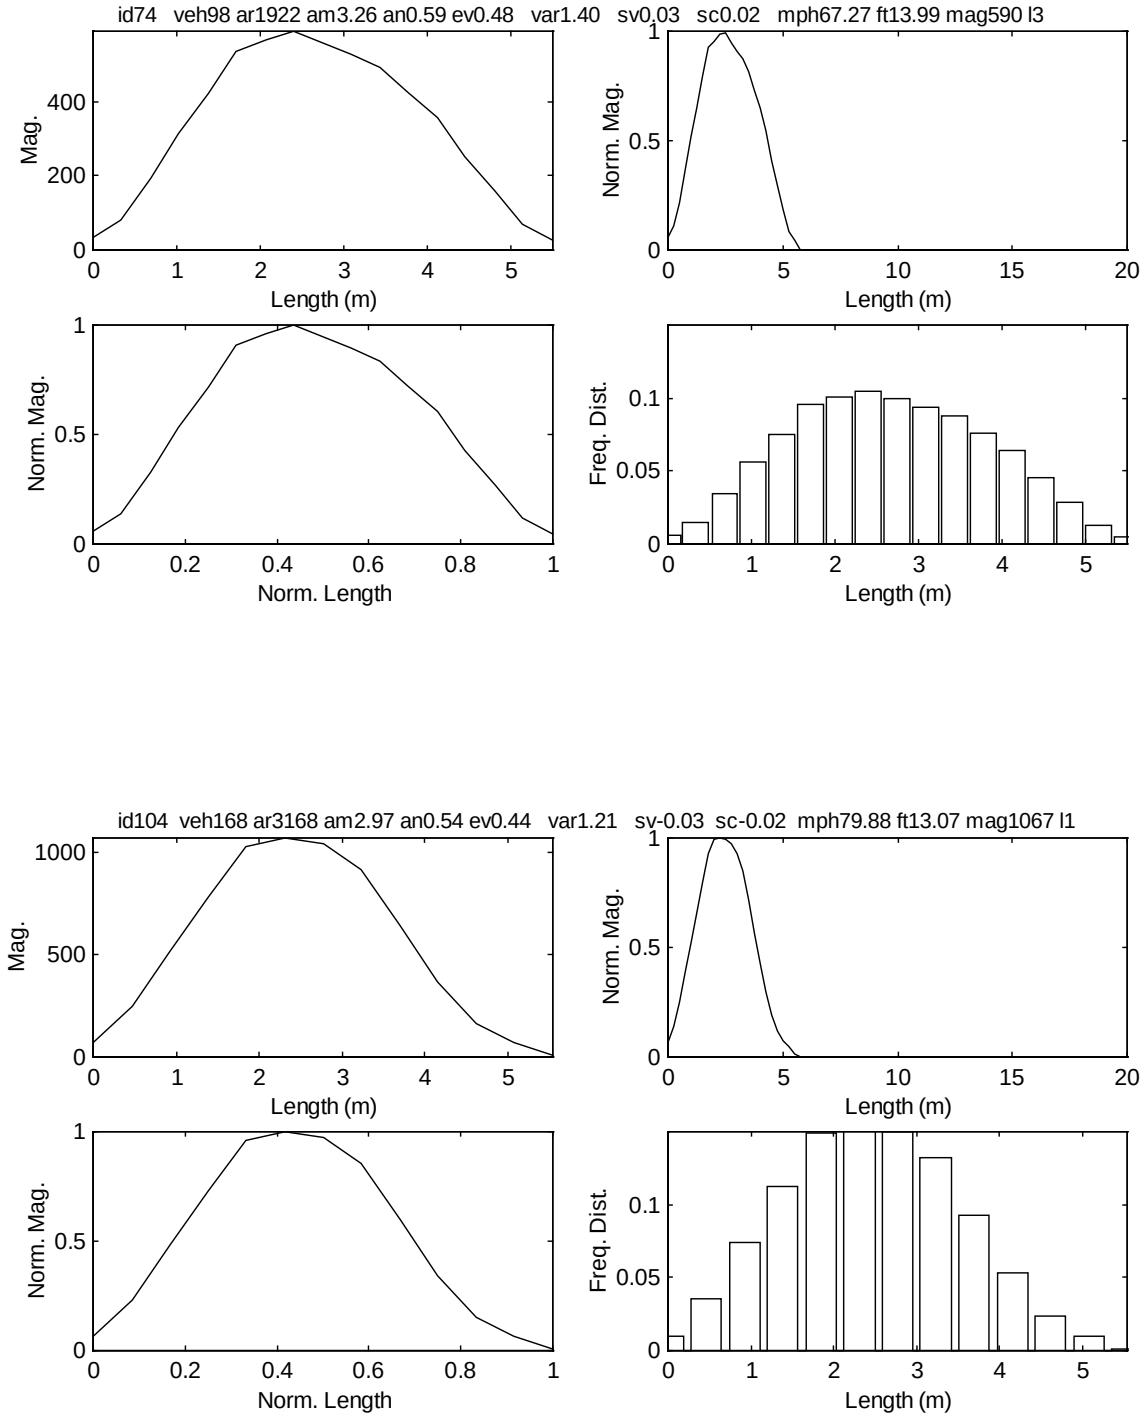

Figure 23. Examples of Misclassified Passenger Cars

After analyzing the error from the first SOFM implementation, the training set was changed by replace the “cargo van” signature. The change in training set also produced a new four category classification scheme shown in Table 14. The weights of the trained network is shown in Figure 24. After changing the training set by one vehicle the reidentification rate increase to 81%. All the misclassifications were again one class misclassifications, meaning that the class was assigned erroneously to a neighboring class. The results reported are similar to Pursula and Pikkarainen (1994). In their study, they reported successful classifications of around 80% by using two-dimensional self-organizing feature maps.

Table 14. New Classification Scheme

Class	Vehicle Types Represented
1	cargo truck, vehicle with trailers, bus, and trucks with >2 axles
2	limousine
3	van, SUV, full-size truck, pickup
4	passenger car, minivan, sports car, station wagon

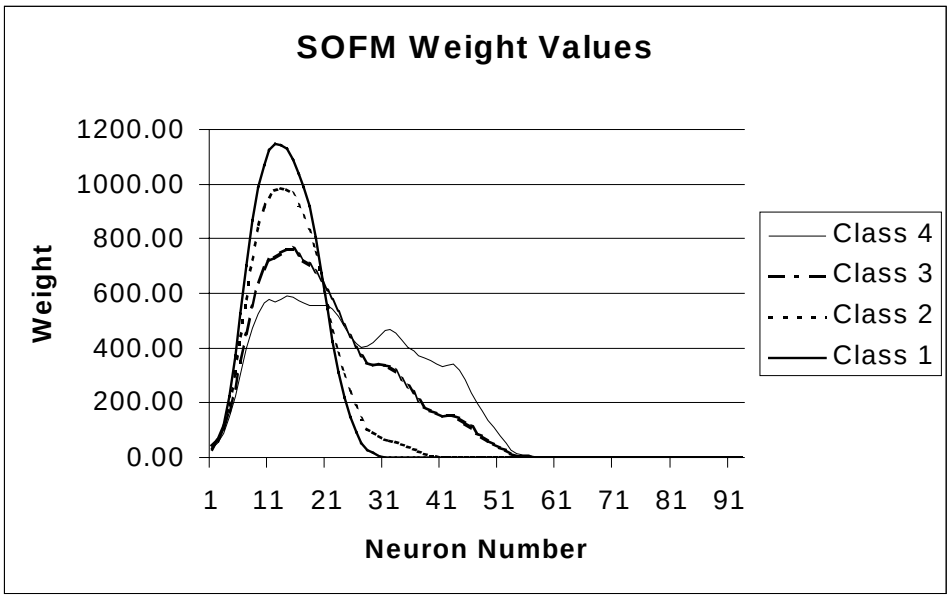

Figure 24. Plot of SOFM Weights

The 4 neuron SOFM was also tested with a second test set composed of 137 vehicle signatures. The classification rate of 80% on the second data set is similar to the results from the first test set.

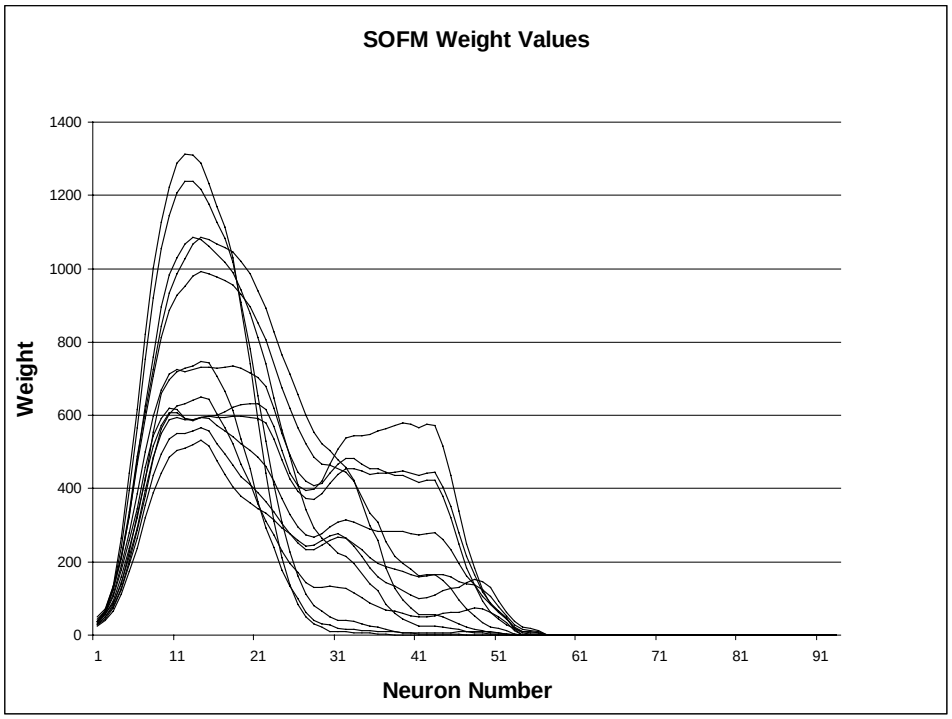

Figure 25. Plot of Neuron Weights for SOFM with 13 Neurons

Table 15. Classification Schemes

Video Vehicle Type	Vehicle Description	Scheme 1	Scheme 2	SOFM Output
1	passenger car	1	1	1,2
2	minivan	1	1	1,2
3	sports car	1	1	1
4	sports utility vehicle	3	2	4,5
5	van	2	2	3
6	limousine	9	7	13
7	2-axle truck	5	3	7,8
8	vehicle with trailer	8	6	12
9	bus	7	5	11
10	full-size pickup	4	2	6
11	pickup	3	2	4,5
12	station wagon	1	1	1
13	> 2-axle truck	6	4	9,10

A second SOFM configuration with 13 neurons was also trained and tested. The resulting weights are plotted on Figure 25. The thirteen outputs were mapped to two test configurations, one with 9 vehicle classes and the other with 7 vehicle classes. The two schemes are listed in Table 15. Column one is a classification based on video. Column 3 lists the mapping for the 9 vehicle class scheme. Column 4 lists the mapping for the 7

vehicle scheme. Column 5 lists the 13 neuron SOFM outputs corresponding to the video classification.

Table 16. SOFM Results

Overall Classification Rates	
9 Vehicle Classes Data Set 1	71%
7 Vehicle Classes Data Set 1	87%
7 Vehicle Classes Data Set 2	82%

The classification rate using 9 vehicles classes was 71% for the first data set. Most of the error is due to SOFM's inability to distinguish between 3 vehicle classes: 2, 3, and 4. In other words, cargo vans, SUV's, pickups, and full-size pickups exhibit very similar signatures. This is not surprising, since smaller SUV's share platforms with pickups, and bigger SUV's share platforms with full-size pickups. By expanding the SUV/pickup cluster to include cargo vans and full-size pickups, a 7 class scheme results. The overall classification rate for that scheme is 87% with 100% accuracy for classes 3 and 7, and 98% accuracy for class 2. The results of the overall classification rates are listed in Table 16.

The overall result of testing using the second data set was 82%. The rate for individual classes were 74% for class 1, 94% for class 2, 88% for class 3, 75% for class 4, and 100% for classes 5 through 7. One major reason for the errors was the misclassification

of off-center passenger cars as SUV's or pickups. As was mentioned before, such vehicles produce a small inductance change and confuse the SOFM.

CHAPTER 6. INDIVIDUAL VEHICLE SPEED ESTIMATION USING SINGLE LOOP INDUCTIVE SIGNATURES

The computation of vehicle speed is an important component of vehicle classification systems since speed is utilized for transforming the raw signatures. Previously, speeds were mainly acquired from double inductive loops configured as speed traps, since single loop speed estimates based on assumptions of a constant vehicle length were inaccurate. However, more accurate measurements of speed can now be accomplished with single loops by utilizing inductive signatures of vehicles that are outputed from newer detector cards. An algorithm using signal processing and statistical methods was developed to extract speeds from inductive signatures. The results show that the proposed algorithm performs better than conventional single loop estimation methods. The results also show that the algorithm is robust under different traffic conditions and is transferrable across surveillance sites without the need for recalibration. The use of the extensive single loop surveillance infrastructure is a cost-effective way of obtaining more accurate network-wide vehicle classification information.

6.1 REVIEW OF SPEED MEASUREMENT METHODS

There are two ways of measuring speed using inductive loop detectors. The first way is to use a set of double loops in a speed trap configuration. The speed trap refers to the measurement of the time that a vehicle requires to travel between two detection points (Woods et al., 1994). Woods et al. also discuss other aspects of the speed trap design including optimum speed trap spacing. Speed is measured as

$$v = \frac{d}{t_{on}^{2} - t_{on}^{1}}$$

Where:

v = speed

d = distance between detectors

t_{on}^{1} = time when the first detector turns on

t_{on}^{2} = time when the second detector turns on

In addition to the turn on times, the Traffic Control Systems Handbook (Wilshire et al., 1985) recommends the use of turn off times also in order to improve the speed accuracy. The

$$v = \frac{1}{2}\left(\frac{d}{t_{on}^{2} - t_{on}^{1}} + \frac{d}{t_{off}^{2} - t_{off}^{1}}\right)$$

averaging of the turn on and off times results in the following equation:

Where:

t_{off}^{1} = time when the first detector turns off

t_{off}^{2} = time when the second detector turns off

$$v = \frac{l}{t_{off} - t_{on}}$$

The second way of measuring speed is by using a single loop detector. In other words,

Where:

l = effective length, vehicle length + detector length

t_{on} = time when the detector turns on

t_{off} = time when the detector turns off

Since actual vehicle lengths are not known, a mean vehicle length must be assumed for this computation. Because lengths can vary greatly from vehicle to vehicle, the use of a mean vehicle length can result in inaccurate measurements of speed. For example, longer trucks will produce underestimated speeds, and shorter passenger cars will produce overestimated speed.

Another approach of estimating average speed using single loops is by using fundamental traffic flow considerations. The average speed is estimated by using average lane volume and loop occupancy. Some publications that discuss this approach include Athol (1965), Hall et al. (1989), and Jacobson et al. (1989). Speed is expressed in these approaches as

$$v = \frac{q}{o * g}$$

Where:

v = space mean speed

q = flow

o = occupancy

g = factor that takes into account the vehicle and detector length

Hall et al. discusses the use of a sliding value of g which changes with occupancy and is calibrated separately for each set of geometric conditions. This approach assumes that traffic flow is uniform, occupancy is a constant multiple of density, and vehicle lengths are constant. Speed inaccuracies result when the assumed conditions are invalid. Mikhalkin et al. (1972) shows that an estimate using this approach will produce a biased estimate of the space mean speed. Also, individual vehicle speeds are not estimated with this approach.

In order to produce an unbiased estimate of individual and mean speed using single loops, Mikhalkin et al. (1972) suggests the use of the minimum mean square error estimator for speed. This approach assumes that random vehicle arrival times are uniformly distributed. This approach also necessitates the estimation of the unbiased mean vehicle length during the relevant time interval.

The use of single loop vehicle signatures for measuring speed has many advantages over previous approaches of speed derivation using loops. First, it can be implemented with existing single loop infrastructure without cutting new loops for speed traps. Second, there is no assumption of a mean vehicle length, so the accuracy of computing speed is not a function of the distribution of vehicle lengths. Third, in contrast to "traffic flow approaches", there is no assumption of traffic uniformity, and individual vehicle speeds in addition to mean speeds can be obtained. Last, there is no distributional assumptions made on vehicle arrival times or vehicle speeds.

6.2 THE USE OF INDUCTIVE SIGNATURES

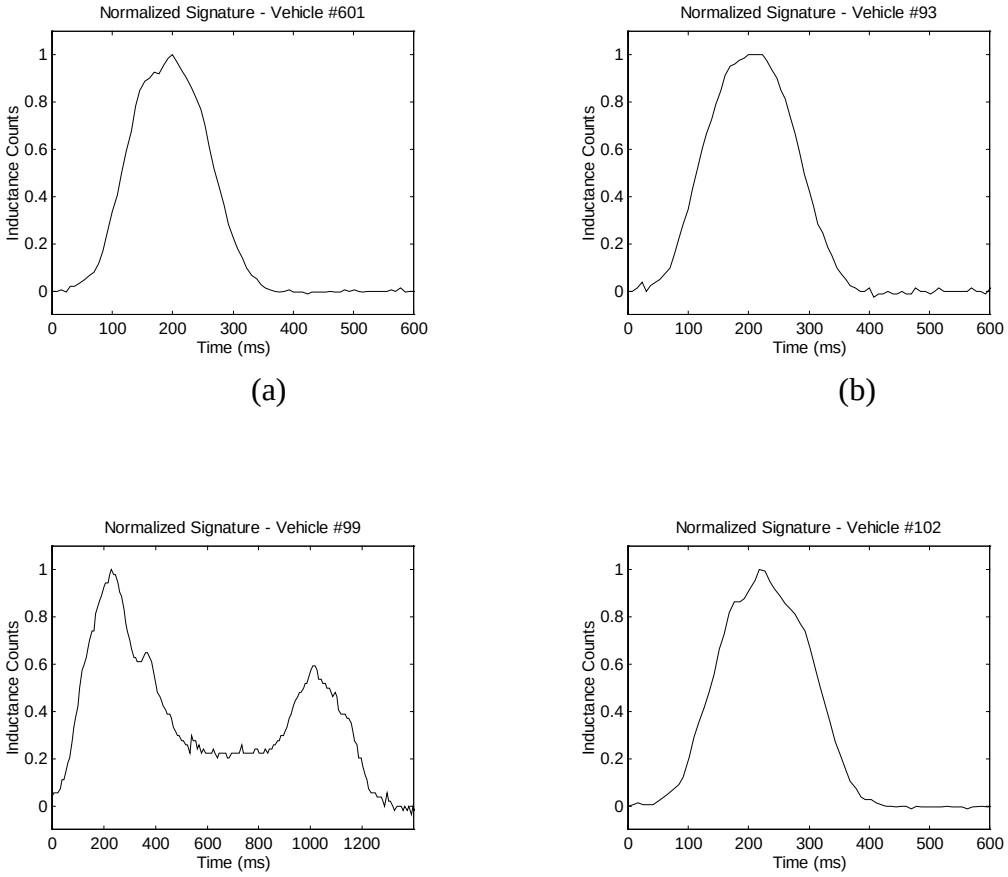

Figure 26. Examples of Vehicle Inductive Signatures: (a) Passenger Car; (b) Sport Utility Vehicle; (c) Semi-Trailer Truck; (d) Minivan

Figure 26 shows four examples of vehicle inductive signatures recorded by sampling the detector at short time intervals. The vertical axis represents the relative inductance change and the horizontal axis represents time. A continuous curve is shown instead of the actual discrete curve in order to clearly display the shapes of the signatures.

As was shown in Figure 26, the vehicle inductive signatures contain both a leading and a trailing edge. These edges are slew rates which represent the rate of the metallic mass of the vehicle moving over the loop magnetic field. The signature within these edges is produced when the entire vehicle is covering the loop. This portion of the signature is no longer the slew rate because of the integration property of the loops, and produces no useful speed information. The estimation of vehicle speed using single inductive loop detectors involves two main procedures. The first is the extraction of the vehicle slew rate or edges from the inductive vehicle signature. The second is the estimation of the vehicle speed based on the slew rate.

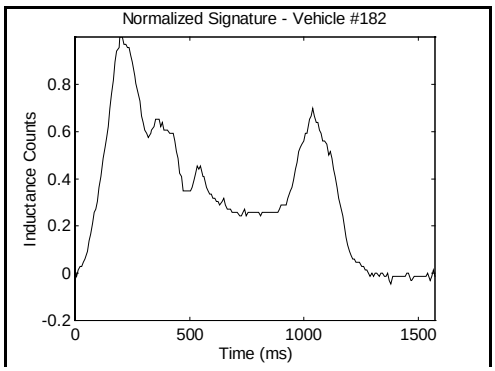

(a)

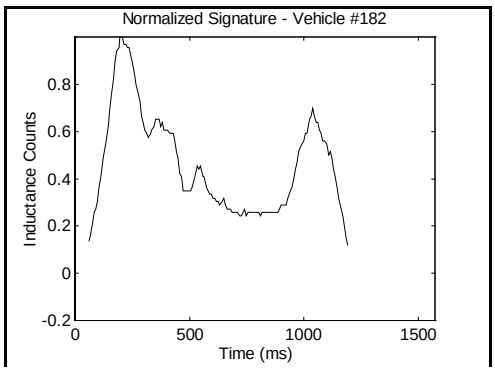

(b)

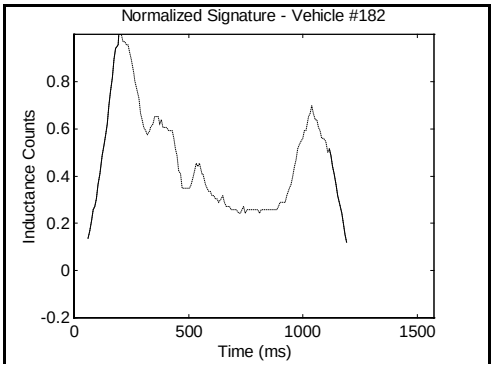

(c)

Figure 27. Slew Rate Extraction Process: (a) Sample Vehicle Signature; (b) Removal of Base Oscillations; (c) Leading and Trailing Edge Slew Rates

There are several steps to extracting the slew rate information from the inductive vehicle signature. Figure 27(a) shows one example of a vehicle signature that will be used to illustrate the process of slew rate extraction. The y-axis has units of percentage change of inductance counts, and the x-axis has units of milliseconds. Note that the vehicle signature which results from decreases in inductance have been inverted to facilitate analysis. The first step in extracting slew rates is to eliminate signature oscillations that have inductance values near the baseline inductance. This is accomplished by using a arbitrary threshold that will eliminate all oscillations while preserving as much signature information as possible. A value of 10%-20% was found to be a good threshold from empirical trials with different datasets. The oscillations near the baseline occur because the detector card is constantly trying to adjust the base inductance value in order to compensate for changes in absolute inductance due to environmental and other effects. The oscillations also occur because vehicles are composed of a complex combination of metallic masses including the undercarriage, the engine, and axles. Figure 27(b) shows an example of a signature with the oscillations removed.

The next step is to extract the leading and trailing edges of the vehicle signature. These edges are slew rates in the sense that they represent the rate of movement of the vehicle's metallic mass over time. The signature edges in this case have units of % inductance change over time. The upper end of the signature edges is found by detecting a local maximum point in the vehicle signature. Therefore the termination point of the signature edge is defined as

$$f(x) \leq f(c), \forall x \in S$$

Where:

c = the time at which the local maximum occurs

$f(c)$ = the local maximum

$f(x)$ = the value of the signature or inductance change at a time x

S = the set of time values that include a single local maximum

Since the signature is actually a discrete signature, the termination point of the leading and trailing edges are found by searching for a local maximum by starting from the threshold value that is illustrated in Figure 27 (b). The resulting leading and trailing edges are shown as solid curves in Figure 27 (c).

There were two data collection periods used in the single loop speed investigation. The morning peak data was collected between approximately 8:00am and 9:30am at the downstream station. This dataset contained 581 vehicle signature pairs. The average flow over that time period was 612 VPH for two lanes. Due to arterial signalization and varying traffic demand, different speeds, acceleration profiles, and traffic flow levels were observed during the 1-1/2 hour period. The maximum observed speed was 30.66m/s (68.68mph). The minimum speed was 5.47m/s (12.25mph). The arithmetic mean speed was 20.77m/s (46.52mph). The standard deviation of the speed was 4.11m/s. The longest electronic length (length of detector+vehicle) observed was 20.31m. The shortest electronic length observed was 3.89m. The mean electronic length was 4.83m.

The midday data was collected between approximately 12:00pm and 1:00pm at the upstream station. This dataset contained 530 vehicle signature pairs. The average flow over the midday period was 575 VPH. The maximum observed speed was 32.52m/s. The minimum speed was 15.16m/s. The arithmetic mean speed was 22.13m/s. The standard deviation of the speed was 2.61m/s.

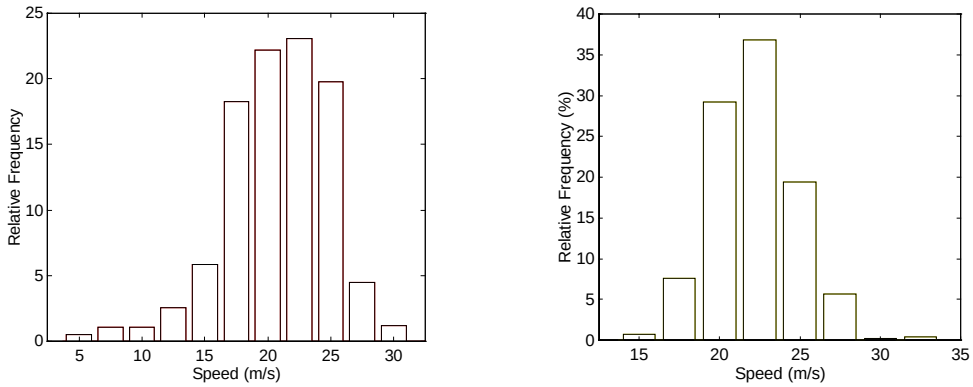

(a) (b)

Figure 28. Speed Distributions (Percent Relative Frequency): (a) AM Peak Downstream; (b) Midday Upstream

Figure 28 shows the percent relative frequency speed distributions for the two datasets. Both the morning and midday datasets contained a wide range of traffic flow conditions and vehicle speeds. This diversity is useful for developing and testing speed computation algorithms. The

two datasets are collected from two different sites at different times of the day. Therefore these datasets can be used to test the transferability of the algorithms.

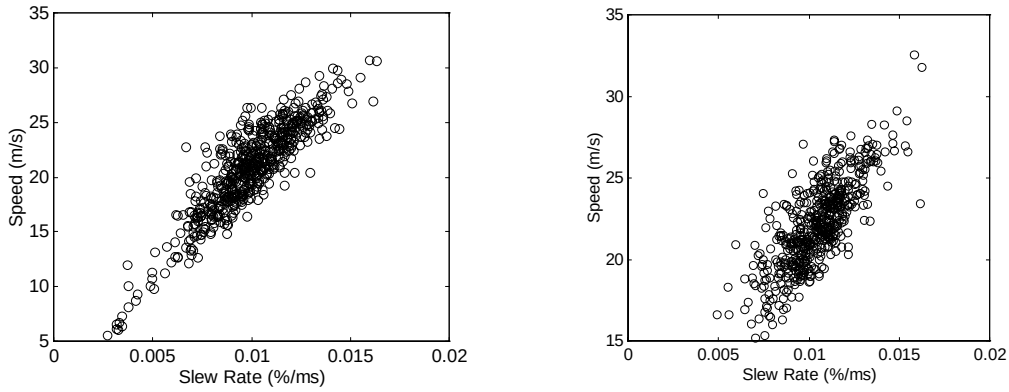

(a) (b)

Figure 29. Scatter Plot of Slew Rate and Speed: (a) AM Peak DownstreamUsing Loop1; (b) Midday Upstream Using Loop 2

The first step in the development of the speed computation algorithms is to plot the available data and to analyze the data graphically. Figure 29 confirms our intuition that the slew rate is linearly correlated with speed. The correlation coefficient for the slew rate and speed is 0.8866 for morning peak data and 0.7762 for midday. Furthermore, Figure 29 also shows that a linear relationship exists between the slew rate and the speed. The scatter plots are also used to compare the data from loop 1 and loop 2 at each double loop station. The plots confirm our intuition that each loop in the double loop station produces data with similar correlation between slew rate and speed. The plots of the morning peak downstream using loop 2 and the midday upstream using loop 1 are not shown as they are almost identical to the plots shown in Figure 29.

The next step is to model the relationship between slew rate and speed in order to predict and/or compute speeds. Given the scatter plots shown in Figure 29, the following simple linear model is postulated for the slew rate/speed relationship of the population.

$$speed_i = \alpha + \beta\ slew_i + \varepsilon_i$$

Where:

$speed_i$	=	dependent variable
$slew_i$	=	independent variable, slew rate
ε_i	=	disturbance term

In the context of linear regression,

$$\hat{s}peed_i = a + b\ slew_i$$
$$e_i = speed_i - \hat{s}peed_i$$

Where:

$\hat{s}peed_i$	=	estimates of the speed
$slew_i$	=	regressor, slew rate
a,b	=	parameters of regression
e_i	=	residual

The following assumptions are made in order to perform linear regression. First, there is the assumption that no observations on $slew_i$ convey information about the expected value of the disturbance. In other words, the conditional expected value of the disturbance is equal to zero. This is a reasonable assumption since only the measurement errors and the random vehicle fluctuations contribute to the disturbance. The random vehicle fluctuations are the result of different driving behavior and vehicle mechanics. Second, the assumption of homoscedasticity is made. Last, the disturbances of the slew rates are assumed to be uncorrelated. The last two assumptions are also reasonable since there is nothing in the loop resonant circuit theory to suggest that variance is not constant across slew rates and that disturbances are correlated. These assumptions are supported by residual tests, and test outcomes are reported in the results section.

Least square estimation, which tries to minimize the sum of square residuals, is used to determine the parameters *a* and *b*. The solution of the following normal equations yields the parameters of the regression model.

$$a\,n + b\sum_{j=1}^{n} slew_j = \sum_{j=1}^{n} speed_j$$

$$a\sum_{j=1}^{n} slew_j + b\sum_{j=1}^{n} slew_j^2 = \sum_{j=1}^{n} slew_j\ speed_j$$

Speed can then be computed for a vehicle by using the linear regression model and a given slew rate for that vehicle.

An interesting experimental concern arises when the standard error of *b* is examined in the following equation (Wannocott, 1990):

$$standard\ error\ \_\ \frac{\sigma}{\sqrt{n}} \bullet \frac{1}{s_{slew}}$$

The above equation implies that by increasing S_{slew}, the standard error of *b* can be reduced. This suggests that one might want to collect the regression data by driving control vehicles over the loops with as wide a spread of slew rates as possible. However, this approach is a costly approach that detracts from the ease of implementation of the proposed speed computation system. Consequently, only data from natural occurring traffic is used for calibrating the proposed system.

6.3 DERIVATION OF GROUND TRUTH

The ground truth for individual vehicle speed is derived by using double loop data. Both digital and analog detector outputs were obtained during the vehicle signature data collection. Double loop speeds that are computed using digital outputs can have typical errors of between 3% and 5% for commonly observed vehicles such as cars and pickups (Woods et al. 1994).

Pursula et al. (1989) found that the standard error of the speeds measured with analog signatures is only one third of the error of the speeds with the traditional digital output from thresholds. The measurement of vehicle speed is determined by using the 50% amplitude point of the leading edge of the signatures from loop 1 and 2 (Pursula et al. 1994). Thus, this technique is used in deriving the ground truth.

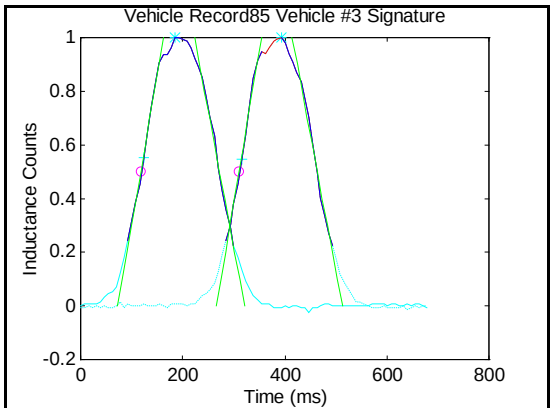

Figure 30. Double Loop Speed Computation Using Signatures

Figure 30 illustrates the double loop speed computation process. In this figure, the first signature corresponds to the first loop in the speed trap, and second signature corresponds to the second loop. For each signature, the front edge is linearized, and the 50% point is determined and shown as a circle in the figure. The vehicle travel time is the difference between the time corresponding to the first circle and the second circle. Since the distance between the double loops is known, the speed is determined by dividing this distance by the travel time.

An alternate method that requires less computation is to use the peak of the signatures. Signature peaks are shown as asterisks in Figure 30. This method is less accurate since a signature can have a plateau instead of a single peak. In fact, the example in Figure 30 exhibits such inaccuracies.

Regardless of which double loop speed computation method is used, there is some error associated with the measured speeds. More precise instrumentation composed of infrared beam curtains or high precision Doppler radars could yield a more accurate ground truth, but the cost of these devices prohibited their use for this research. Therefore the inaccuracy of the ground truth could have contributed to the error in the research results.

6.4 RESULTS

First the results from the linear regression will be shown. Then comparisons will be made with other single loop estimation methods. The sensitivity of the regression parameter with respect to the calibration size will be analyzed, and other model form possibilities will be discussed. In the results section, the term calibration is used to describe the process of estimating regression parameters.

Table 17. Linear Regression Results

Output (1)	Coefficient (2)	Standard Error (3)	t-Statistic (4)
a (intercept)	3.597205	0.461335	7.797389
b (slope)	1739.639	46.04257	37.78327
R^2	0.8273		
Standard Error of regression	1.9028 m/s		

The downstream morning peak dataset was divided into 300 vehicles for calibration, and the remaining 281 for testing. Linear regression was performed by using the first 300 vehicle signatures from the second loop only, so an additional 300 vehicle signatures from the first loop are also available for testing. The regression coefficient as well as several measures that are commonly included in regression analysis are listed in Table 17. The t-statistic shows that all coefficients are significant. The R^2 value confirms that a large fraction of the variance of the speed can be explained by the slew rate. The standard error shows that about two-thirds of the residuals will lie between -1.90m/s and +1.90m/s. The large value for the slope is due to the fact that the units of the independent variable is in term of % inductance change per millisecond.

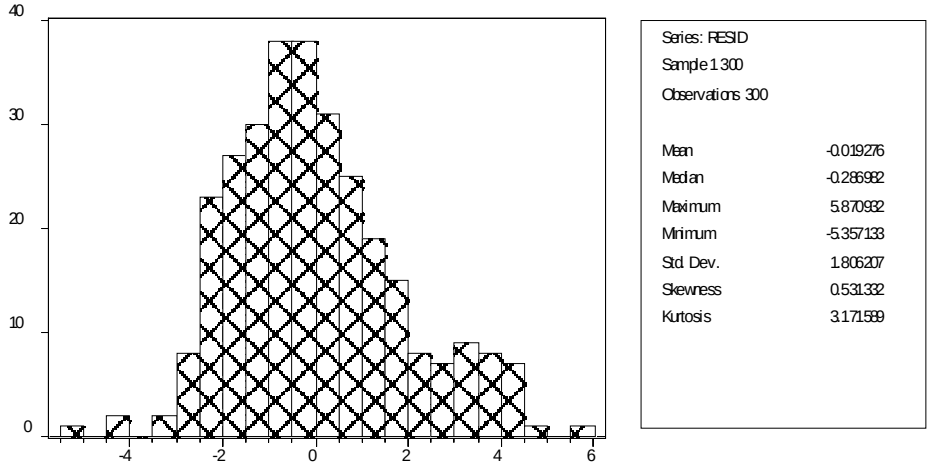

Figure 31. Regression Residual Analysis

Some analysis was performed on the regression residuals in order to validate the normality assumptions of the error term (ε) and to verify that the error terms are not correlated. The error term is assumed to be distributed $N(0,\sigma_\varepsilon^2)$, which means that it is normally distributed with zero mean and constant variance. Figure 31 shows the histogram of the residuals along with the table of descriptive statistics. The histogram and the statistics both show that the distribution of the residuals do not deviate significantly from normal, and the expected value of the distribution is almost zero. In order to test for heteroscedasticity, hypothesis testing is performed using the following null and alternate hypotheses:

$$H_o : \sigma_i^2 = \sigma^2, \forall i$$
$$H_1 : any\ \sigma_i^2 \neq \sigma^2$$

The application of White's heteroscedasticity test (White 1980) produced a small F-value of 0.6172, which means the null hypothesis cannot be rejected and the errors are homoscedastic. The Durbin-Watson statistic was 1.655 which shows only small evidence for autocorrelation, thus the model is not adjusted for time-ordered effects.

Table 18. Single Loop Speed Computation Results Using AM DownstreamData

Errors (1)	Signature Speed (2)	Unbiased Speed (3)	Biased Speed (4)
Average Error (%)	6.7	12.7	36.9
Standard Deviation of Error (%)	5.3	8.7	11.8

Table 18 shows the results comparing three different single loop speed computation methods. The test data is composed of 281 vehicles from the downstream morning peak data. The first column, labeled signature speed, is the speed computed by the approach of using inductive signature slew rates. The second column, labeled unbiased speed, is the speed computed in the conventional fashion by using an unbiased estimate of the mean effective length (length of vehicle and detector). The estimated effective length was found to be 4.83 meters for the downstream morning dataset. This length is divided by the loop traversal time of the vehicle to obtain unbiased individual vehicle speed. The third column, labeled biased speed, is the speed computed by using a commonly assumed effective length of 6.55m (21.5ft) (Arendonk, 1996). An assumed effective length is commonly used by transportation agencies, since it does not require the sampling of vehicle lengths for estimating a sample mean. The data from the second loop of the speed trap is used for testing all three speed computation methods.

The results from Table 18 show that the signature speed algorithm performed better than other conventional speed computation methods. The signature speed was more accurate and shows less variability than the other two approaches. A value of 8.7% standard deviation for the unbiased speed computation error is not surprising, since this standard deviation is a function of the distribution of vehicle lengths. In locations where the traffic is more homogeneous in term of vehicle type for all times of the day, the unbiased speeds is expected to be much more accurate and to exhibit less variability than the results from Table 18. The biased speed shows a

significant average error of 36.9%, since a biased effective length is assumed. This result shows that if an effective length that does not reflect the traffic characteristics of the particular location is assumed, then significant speed computation error will result.

Table 19. Testing Transferability Using AM Downstream Data

Errors (1)	Signature Speed Using Second Loop (calibrated) (1)	Signature Speed Using First Loop (uncalibrated) (2)
Average Error (%)	6.7	7.2
Standard Deviation of Error (%)	5.3	5.9

In order to test the transferability of the calibrated algorithm and regression coefficient, speeds are recalculated using the vehicle signatures from the first loop instead of the second loop. In other words, the second loop signatures are used for calibration, while the first loop signatures are used for testing. The results in Table 19 show the average error and the standard deviation of

error increased only slightly when using testing data. Table 19 presents some evidence for the transferability of the calibrated algorithm.

Table 20. Single Loop Speed Computation Results Using Midday Upstream Data

Errors (1)	Signature Speed (2)	Unbiased Speed (3)	Biased Speed (4)
Average Error (%)	6.7	7.24	37.1
Standard Deviation of Error (%)	5.7	11.73	15.9

In order to further verify the transferability of the calibrated algorithm, testing is performed using the vehicle signatures from a different location recorded at a different time of the day. This dataset is recorded at the upstream site during midday. The results in Table 20 reaffirm the robustness of the algorithm. The average error remained the same while the standard deviation of error increased only slightly. It is interesting to note that the unbiased speed, computed using

the unbiased mean effective length of this dataset, shows less error than the morning dataset. This result reflects the more homogeneous midday traffic that contains less truck traffic and more lunch time passenger vehicles.

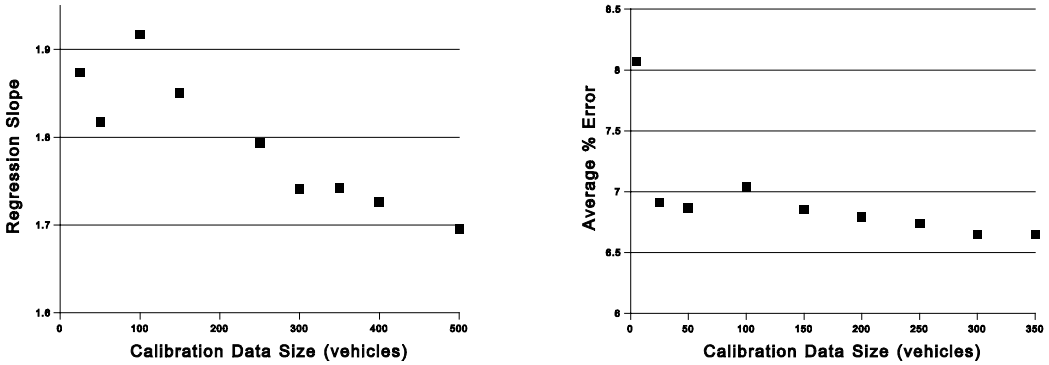

(a) (b)

Figure 32. Sensitivity to Calibration Data Size: (a) Sensitivity of Regression Parameters; (b) Sensitivity of Average Prediction Error

The sensitivity of the regression parameter with respect to the calibration size and the associated average estimation errors are shown in Figure 32. Figure 32 (a) shows that the regression slope fluctuates when less data is used for calibration, but the value levels off after 300 vehicles are used. This is the reason for using a calibration dataset size of 300 vehicles. Figure 32 (b) shows the sensitivity of the average prediction error using a test data size of 231 vehicles. In terms of average prediction error, errors of around 7% are obtained after the calibration size is larger than 25 vehicles. However, regression theory shows that a larger variability in the independent variable leads to smaller standard errors in the residuals; therefore, a more conservative

calibration size such as 300 vehicles will most likely contain more variability. Regression can also be accomplished by selecting a small set of vehicles that have a wide range of speeds, but such a set is not contiguous in terms of arrival times and does not lead to real-time calibration.

Table 21. Regression Using Other Model Forms

Output (1)	Coefficient (2)	Standard Error (3)	t-Statistic (4)
(a) Quadratic Model Form: speed = a + b slew + c slew^2			
a	-1.838342	1.016288	-1.808879
b	3028.725	221.9280	13.64733
c	-70812.11	11953.21	-5.924108
R^2	0.845554		
Standard Error of	1.802480 m/s		

regression	

(b) Cubic Model Form: speed = a + b slew + c slew^2 + d slew^3

a	-2.564159	2.046421	-1.252997
b	3327.396	783.6314	4.357333
c	-106996.3	89315.42	-1.19960
d	1343666	3286728	0.408816
Standard Error of regression	1.805012 m/s		
R^2	0.84541		

Table 21 shows the regression results of using other models forms to describe the relationship between slew rate and speed. The t-statistic shows that some coefficients of these more complicated model forms were not statistically significant. The R^2 values and the standard errors

of regression are all similar between these model forms and the linear model. Also Figure 29 presents a strong graphical evidence for the use of linear regression. Therefore the simpler linear model is adequate for predicting vehicle speeds.

This chapter presents a new methodology for computing vehicle speeds using single loops. This allows the implementation of vehicle classification systems with only single loops. Single loop speed methodology involves the use of signature slew rates, and a linear regression model is used for deriving speeds from slew rates. Results show that this method can perform better than other methods of computing speed. One main advantage this method has over other methods is the fact that the accuracy of predicting speeds is not a function of the distribution of vehicle lengths. Another advantage, is the robustness of the method that is shown to be temporally and spatially transferable. This method requires very little computing power and cost to implement. The simplicity of the methodology leads to uncomplicated real-time implementations for traffic management and control purposes.

The utilization of the current single loop infrastructure avoids costly road closures and equipment associated with the cutting of double loops. In many cities and states, this infrastructure is extensive and can produce speed and vehicle classification information for different Intelligent Transportation System needs. However, double loops or other detection systems are still necessary if high accuracy in speed computation is required.

Local accelerations can be computed by using this method if double loops are available. The local accelerations are simply the difference between the speed computed from the first and second loops in a speed trap. The local accelerations can be a valuable input for congestion monitoring and incident detection algorithms. No acceleration results were presented because there was no ground truth available for validating results.

CHAPTER 7. CONCLUSION

The vehicle classification algorithms described in this report exploit the current inductive loop infrastructure by collecting inductive signatures. By using these signatures, surveillance capabilities are increased, one of which is the ability to derive vehicle classification data. These data can be used as inputs to various ITS strategies in order to improve the efficiency, safety, environmental sustainability, and maintainability of transportation networks.

The decision theoretic algorithms described in this report presented three different heuristics that combined different feature vectors in different ways to separate vehicle classes. These algorithms are examples from which other heuristic algorithms could be developed for particular ITS applications.

One advantage of using heuristic classification algorithms is the reduction of data transmission and storage requirements that results from using feature vectors instead of raw signatures. Another advantage is the sequential discrimination of the heuristic algorithm that reduces the feature space at each stage of discrimination. This staged procedure also has the benefit of allowing the fine-tuning of each stage (discriminant function) independently. As a result of the aforementioned advantages, these heuristic algorithms are suitable for real-time implementation.

This report also formulated a multi-objective optimization problem for training the heuristic algorithms. The training involves the determination of optimum discriminant bounds using golden section search. The three heuristic algorithms produced encouraging results of 81%-91% overall classification rates. In addition, the individual classification rates were fairly uniform particularly for the third heuristic algorithm.

Another vehicle classification algorithm described in this report involves the use of Self-Organizing Feature Map (SOFM). The SOFM uses the concept of self-clusterization to produce groups of inductive signatures that are similar. The results of this research show that the clusters produced by the SOFM correspond to useful vehicle classes that were validated with video. A very significant advantage of using the SOFM is the small training set of 26 signatures that was sufficient to produce classification rates greater than 80%. Another advantage of using the SOFM is the fact that the neuron weights become similar to the inputs and can be analyzed graphically.

The results from the classification algorithms demonstrate the feasibility of deriving vehicle classification information from inductive signatures. The use of the current surveillance infrastructure coupled with the vehicle classification algorithms has great potential in ITS applications.

Because the classification schemes used by previous researchers differ, it is difficult to compare classification algorithms based on the classification rates alone. In a sense, different technologies yield classification schemes that are more suitable for the particular kind of signal that is detected. Consequently, it might be fruitful in the future to consider a combination of different technologies in order to optimize the classification for a particular transportation application. This becomes increasingly feasible as costs of detector technologies drop in response to demand and innovation.

In addition, a speed measurement system was devised to derive speeds from single loop signatures. Vehicle speed is an important component of the vehicle classification system for it makes the signatures speed invariant. The speed measurement system used the linear relationship between slew rate and speed to derive single loop speeds that were more accurate than previous single loop speed estimation methods.

APPENDIX A

ANALYSIS OF FIFTEEN REPRESENTATIVE VEHICLES

- Feature vectors and video ground truth
  - Plot of raw waveform
  - Plot of waveform with normalized magnitude with vehicle length
  - Plot of waveform with normalized x and y axes
  - Plot of waveform as frequency distribution
  - Video image
  - Plots of the Discrete Fourier Transform of Vehicle Waveforms in Complex Plane
- Table 1. First 14 Sorted Magnitudes of Discrete Fourier Transform Coefficients
- Table 2. First 12 Sorted Angles of Discrete Fourier Transform Coefficients
- Table 3. Karhunen-Loeve Eigen values (first 9 components)
- Plot of zeroeth, first, and second derivatives

id1 veh9 ar3153 am5.94 an0.55 ev0.46 var6.77 sv4.06 sc0.23 mph60.52 ft29.04 mag530 l3

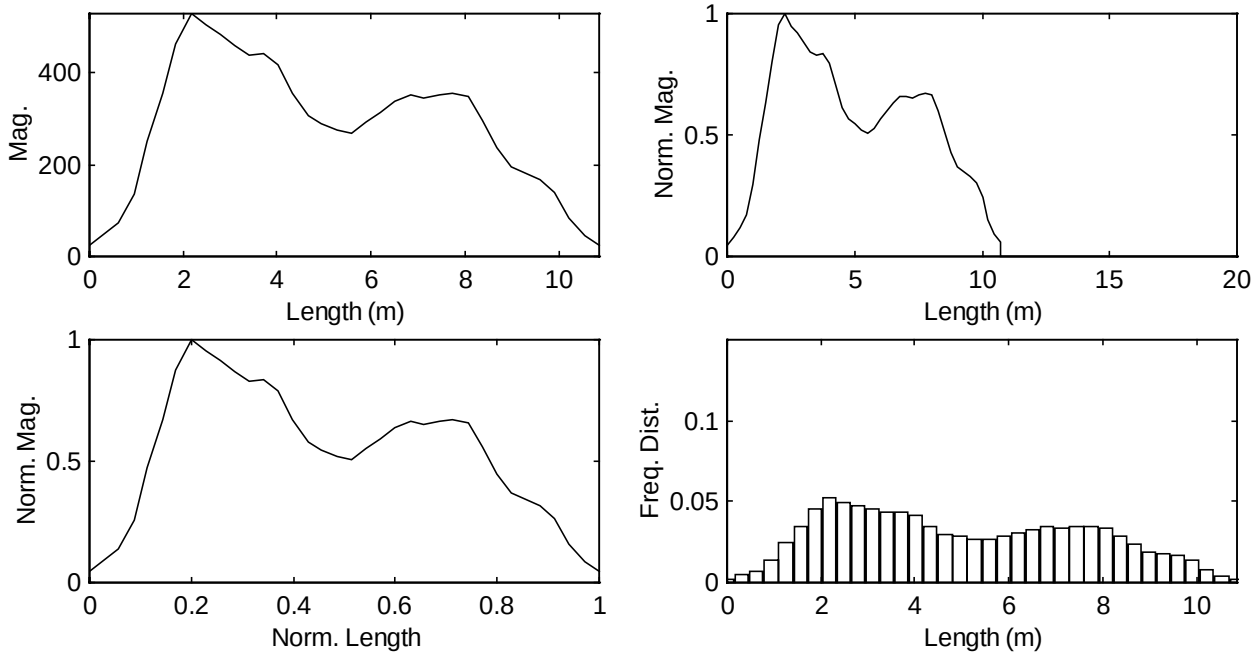

- 

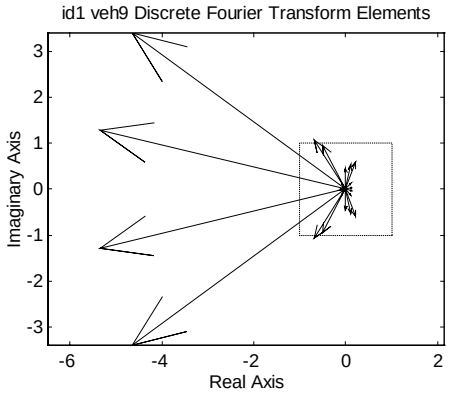

id2 veh14 ar2865 am3.33 an0.52 ev0.49 var1.51 sv-0.07 sc-0.04 mph62.78 ft13.96 mag860 l3

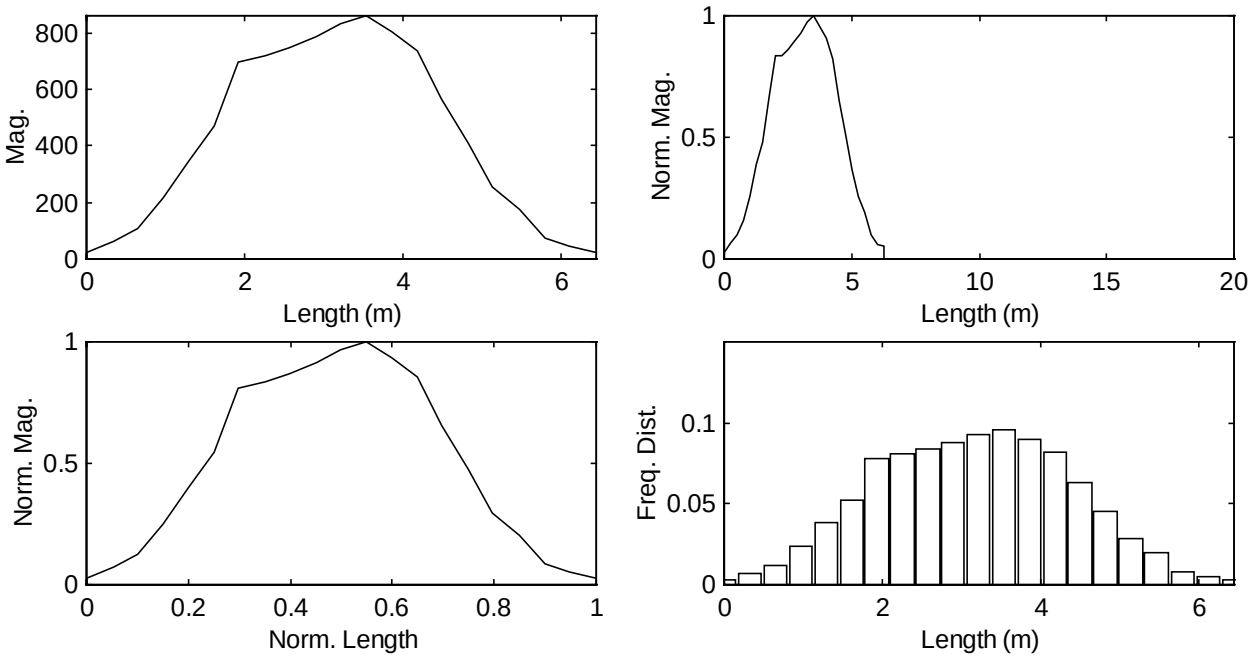

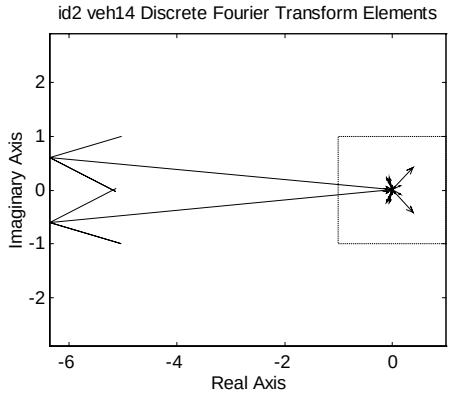

id3 veh27 ar4504 am3.20 an0.50 ev0.45 var1.37 sv0.12 sc0.07 mph65.70 ft13.76 mag1406 l2

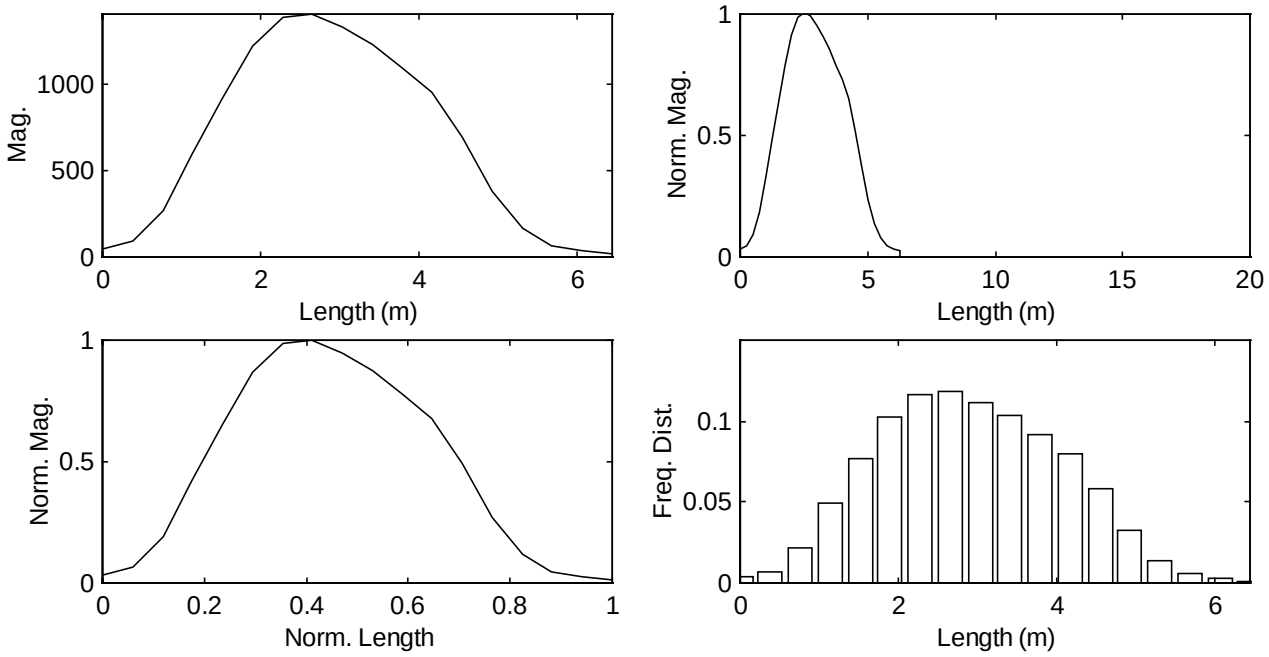

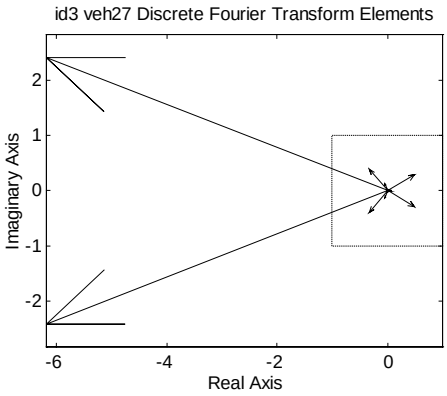

id4 veh28 ar7832 am6.36 an0.64 ev0.49 var4.55 sv0.30 sc0.03 mph59.40 ft25.31 mag1230 l3

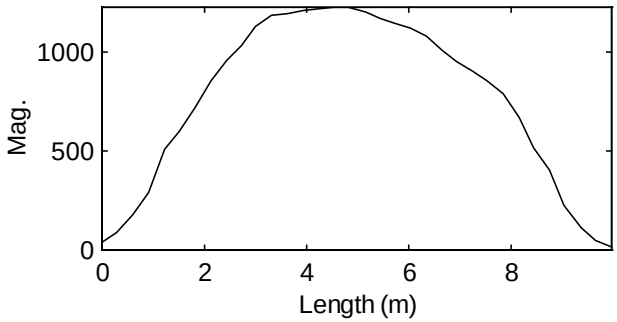

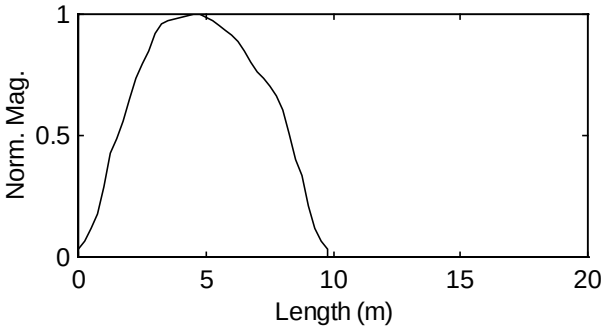

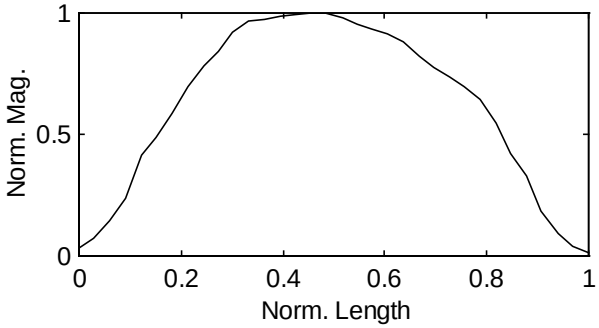

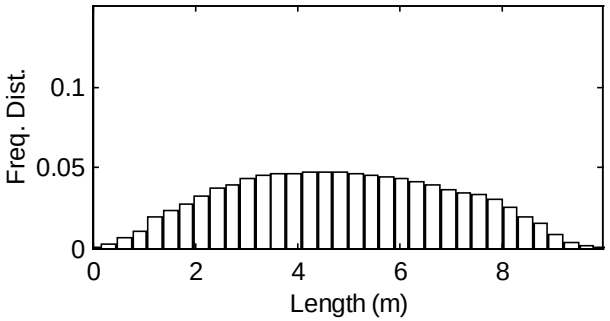

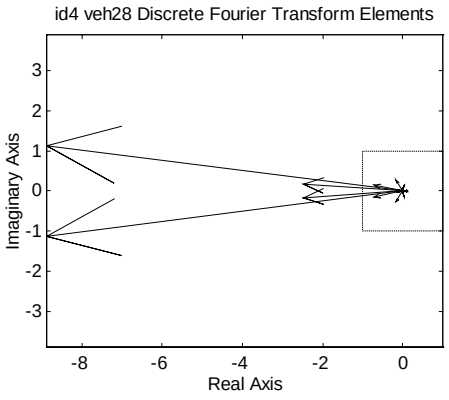

id5 veh34 ar4649 am3.64 an0.56 ev0.46 var1.75 sv-0.06 sc-0.02 mph75.15 ft15.41 mag1277 l1

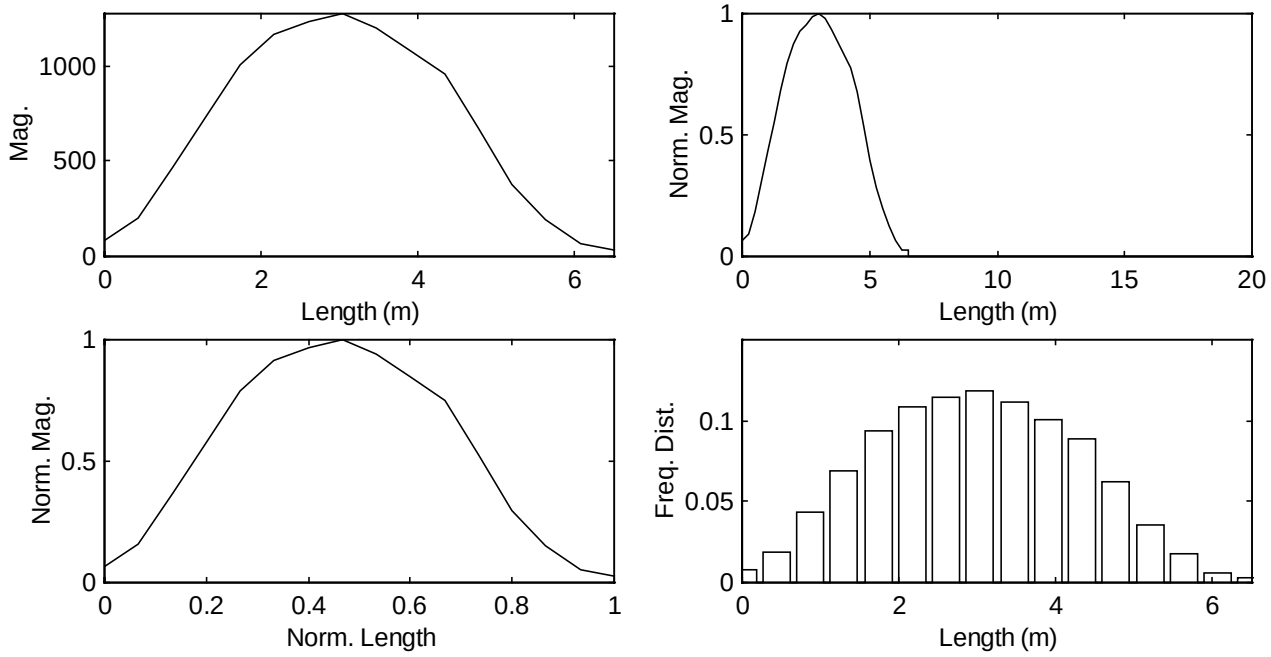

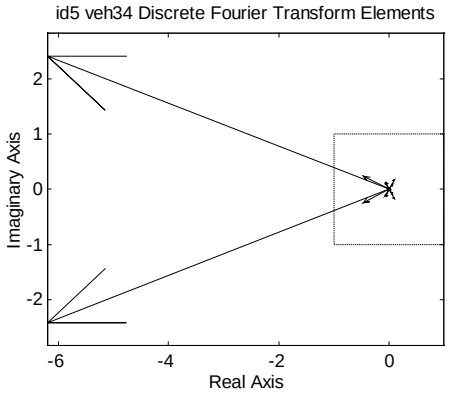

id6 veh35 ar2051 am3.39 an0.52 ev0.50 var1.65 sv-0.02 sc-0.01 mph70.88 ft14.55 mag605 l2

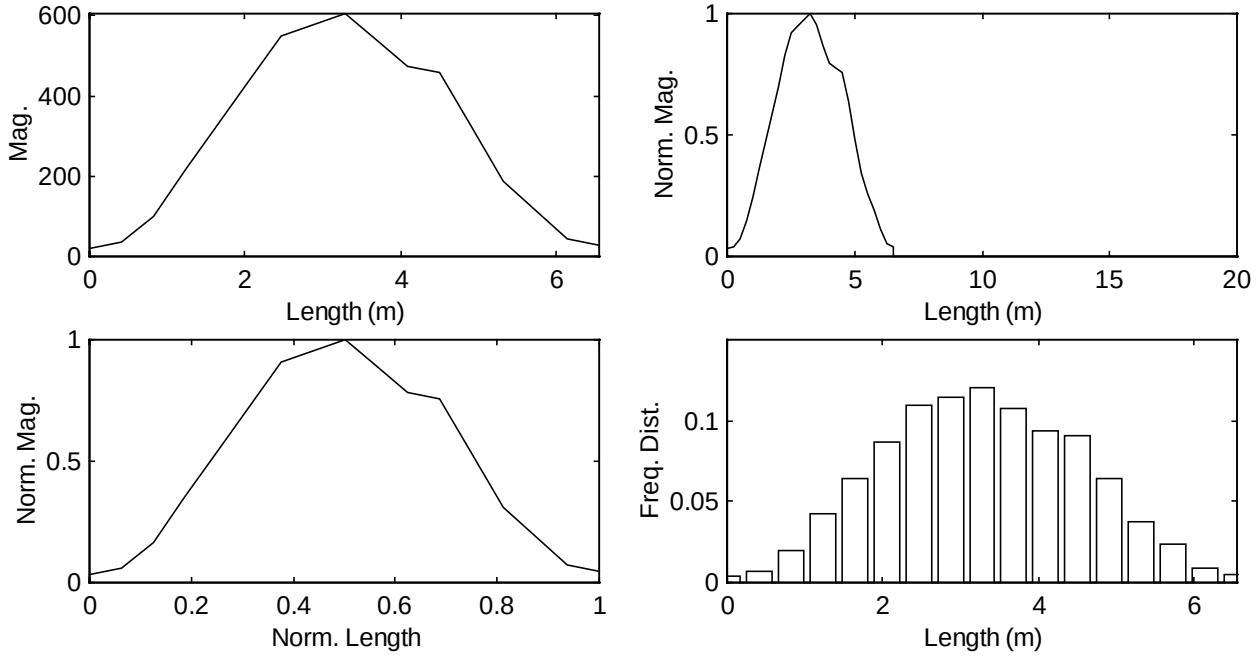

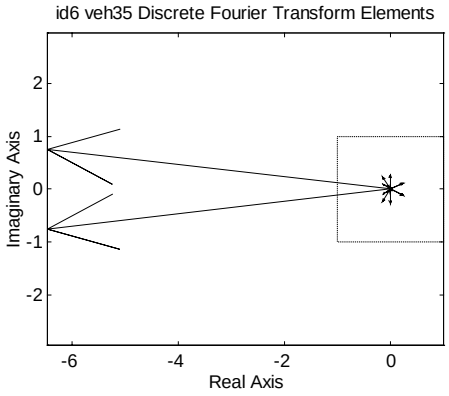

id7 veh90 ar3522 am5.68 an0.53 ev0.44 var5.27 sv2.02 sc0.17 mph60.30 ft27.29 mag620 l4

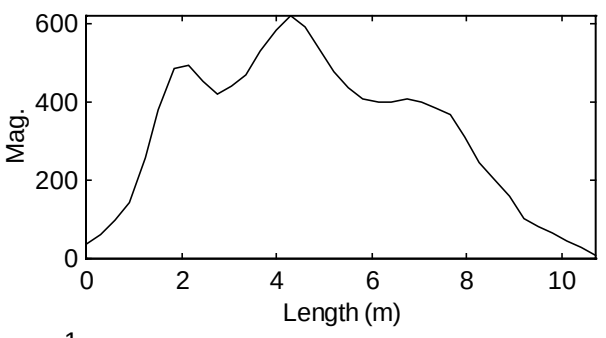

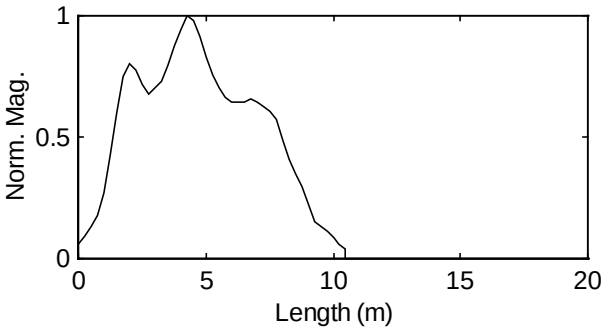

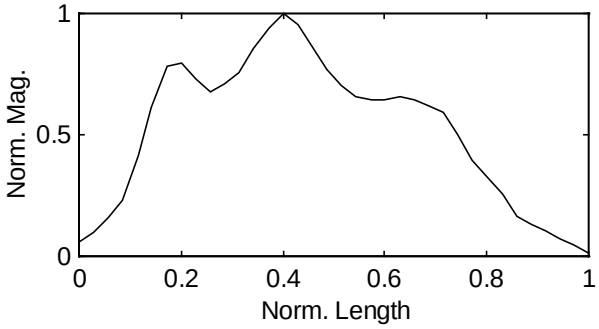

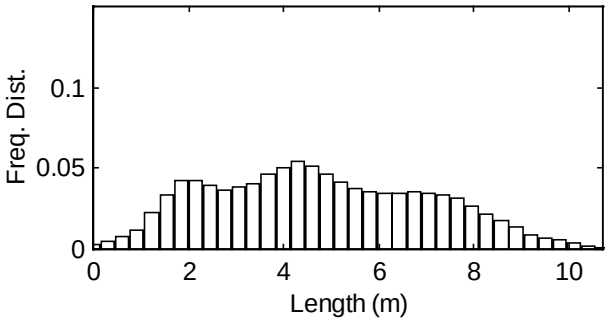

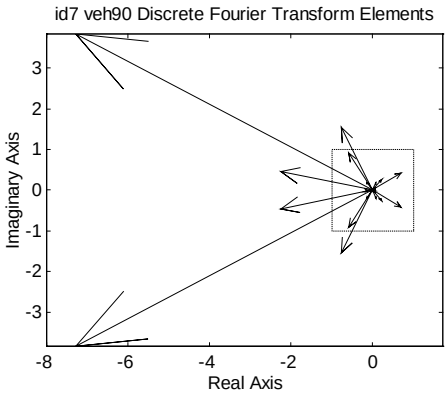

id8 veh104 ar6412 am6.03 an0.29 ev0.39 var28.15 sv114.48 sc0.77 mph63.90 ft57.29 mag1064 l3

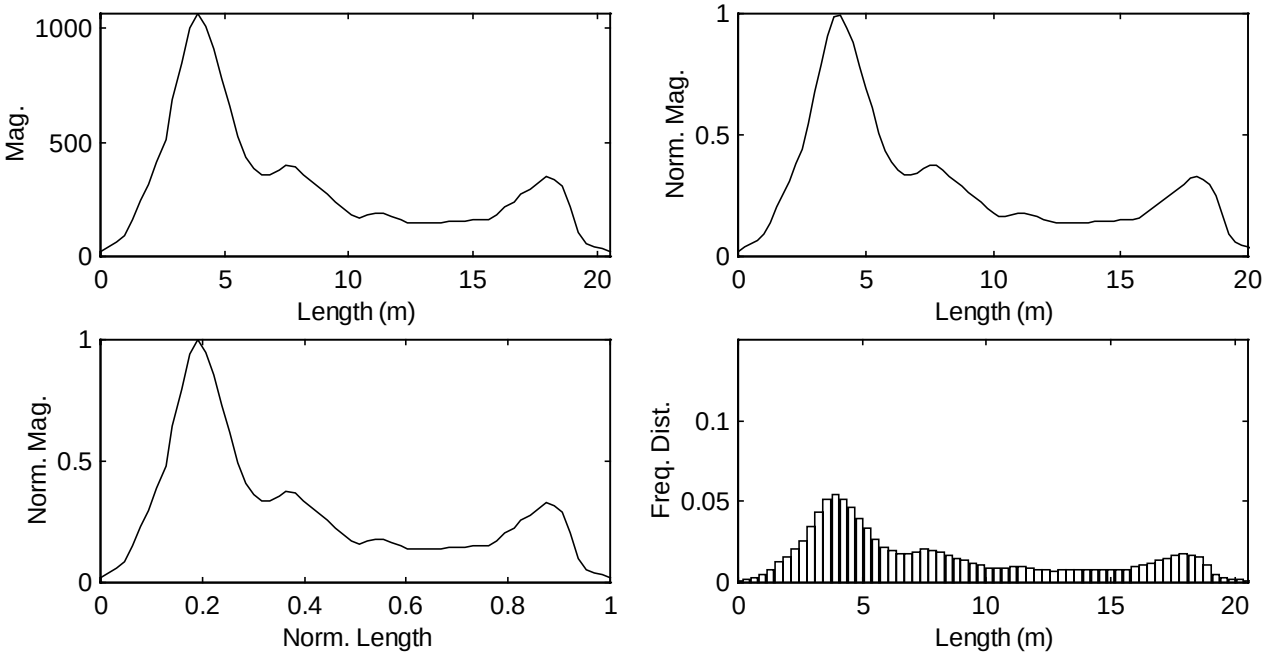

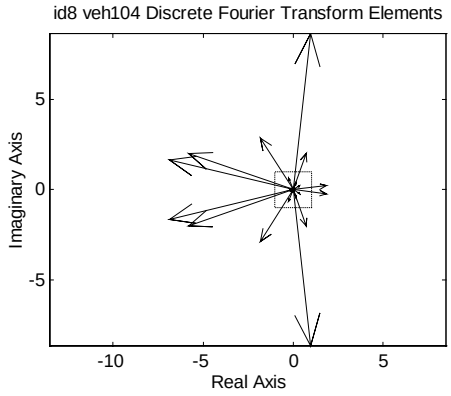

id9 veh117 ar3235 am4.98 an0.60 ev0.46 var2.89 sv0.15 sc0.03 mph67.95 ft20.49 mag649 l3

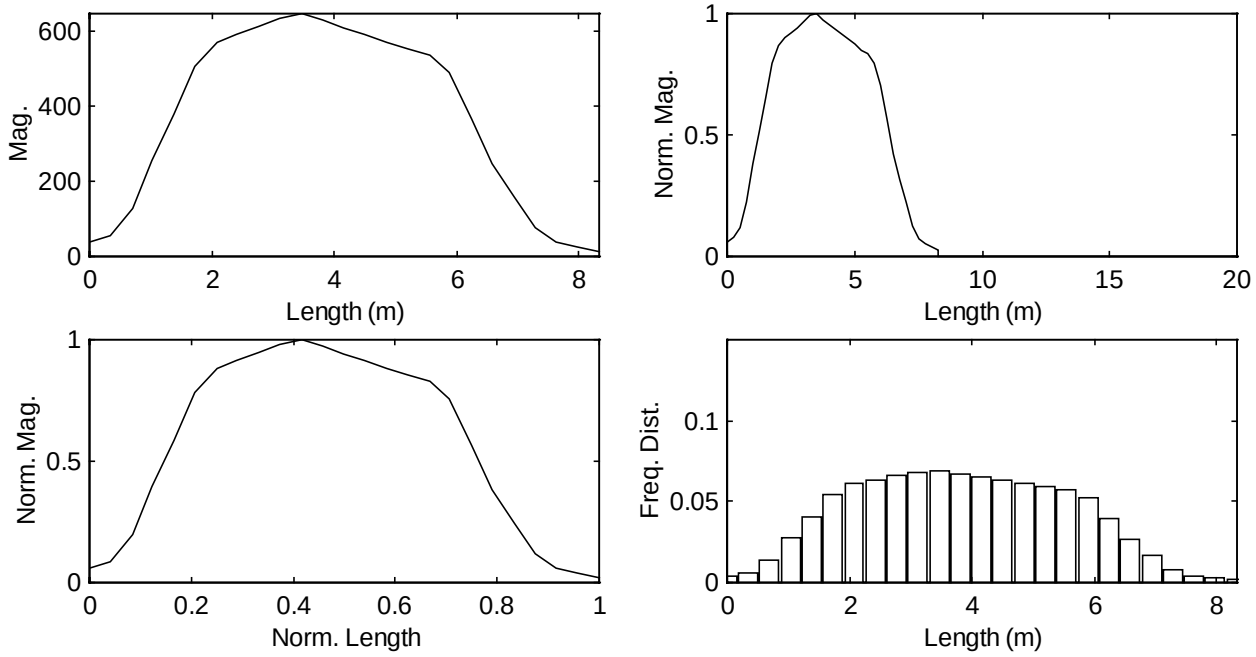

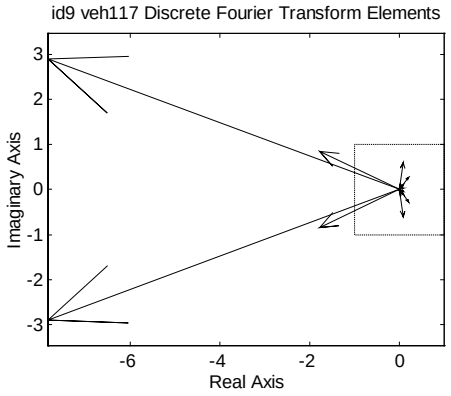

id10 veh124 ar4820 am3.57 an0.51 ev0.47 var1.67 sv0.17 sc0.08 mph62.78 ft15.31 mag1349 l4

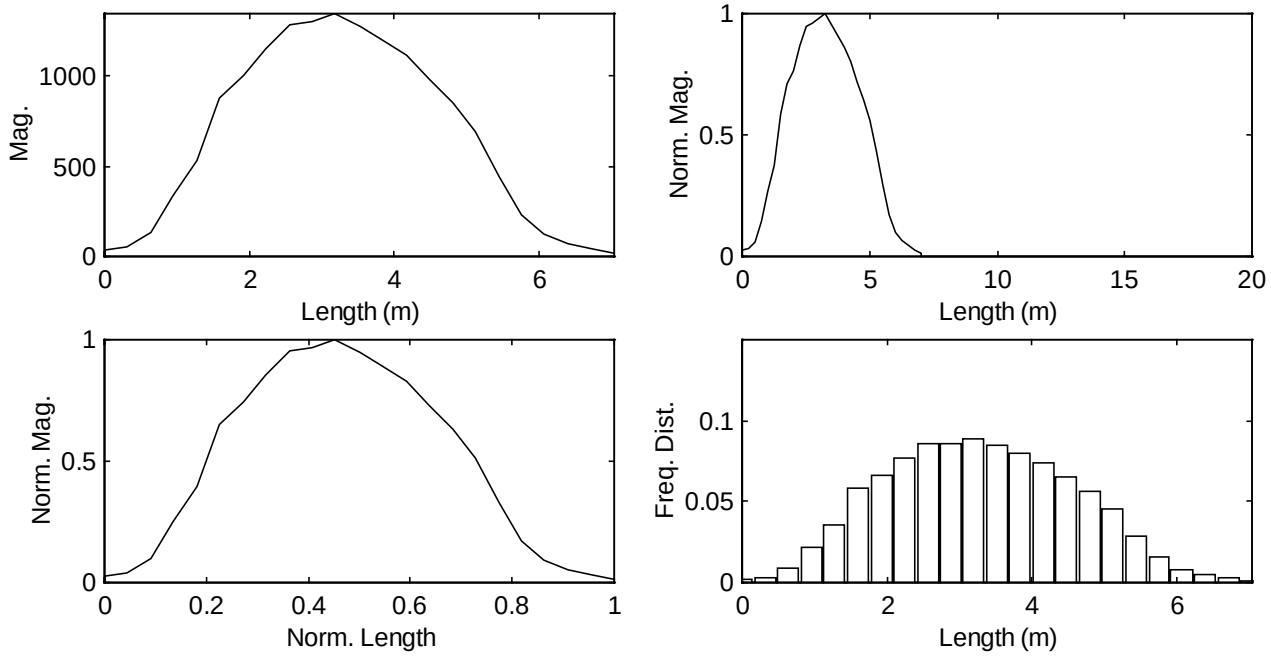

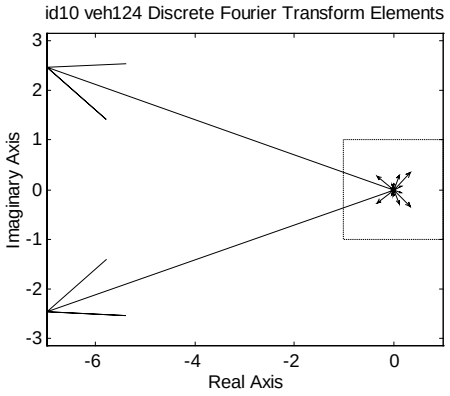

id11 veh166 ar3789 am6.50 an0.47 ev0.44 var13.11 sv20.08 sc0.42 mph61.65 ft40.03 mag583 l3

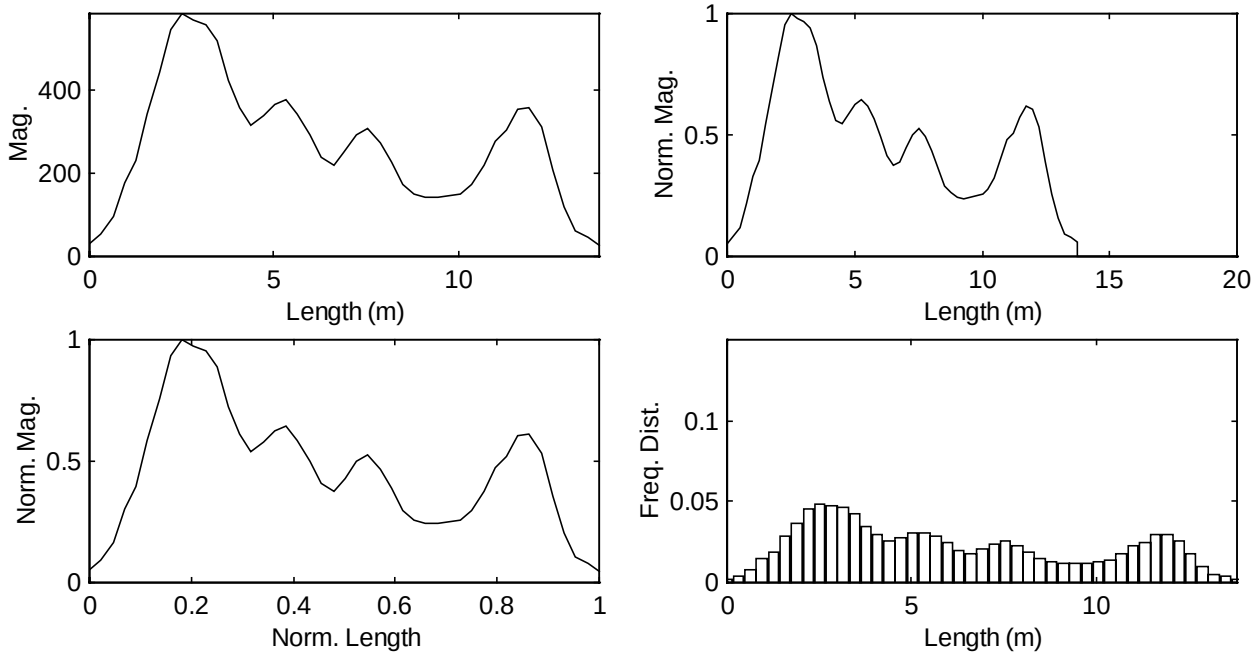

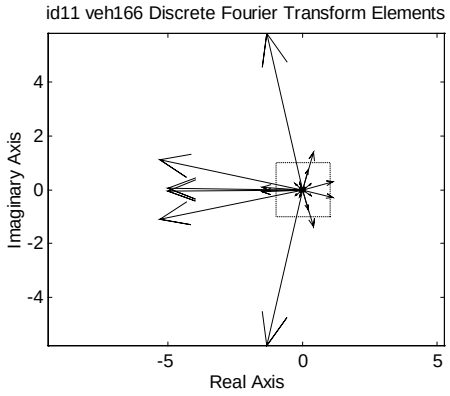

id12 veh167 ar2059 am3.29 an0.53 ev0.47 var1.50 sv-0.12 sc-0.07 mph71.10 ft14.59 mag626 l2

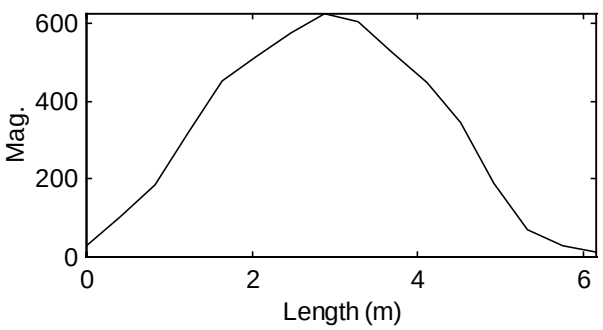

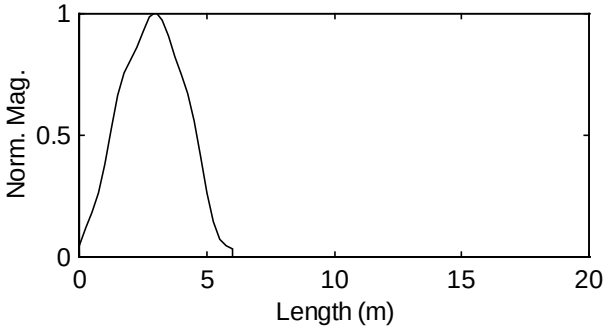

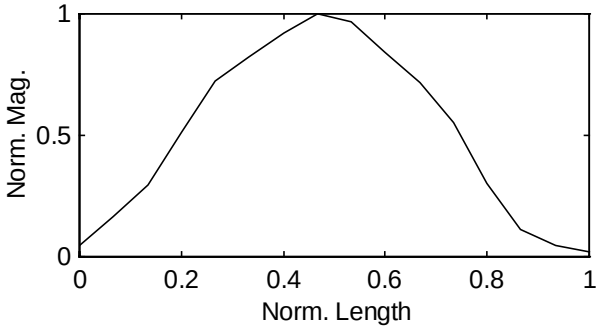

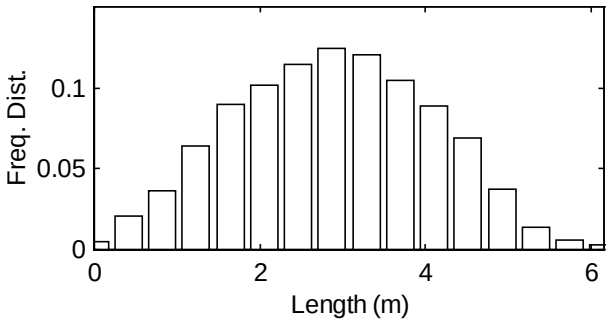

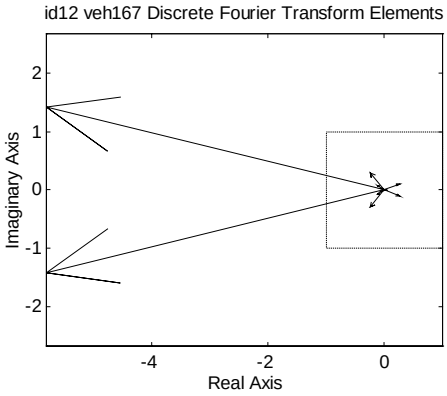

id13 veh321 ar5379 am8.63 an0.45 ev0.45 var28.11 sv36.62 sc0.25 mph59.18 ft55.51 mag623 l3

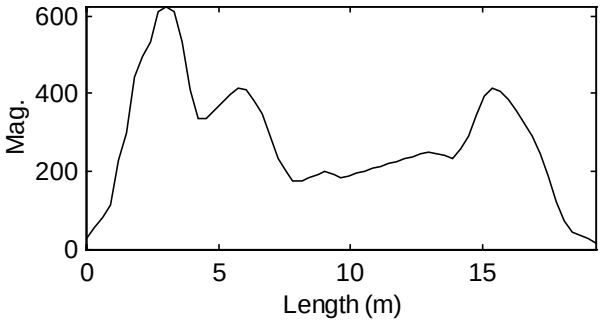

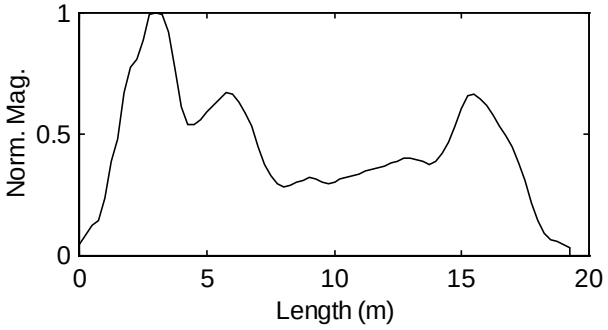

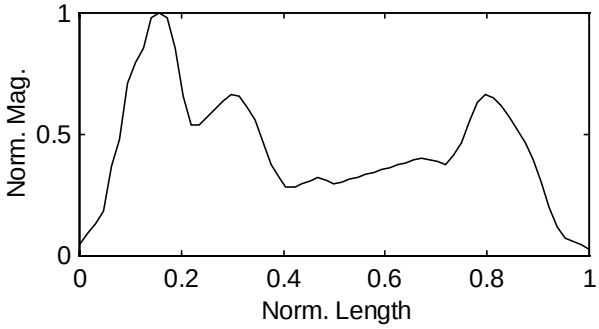

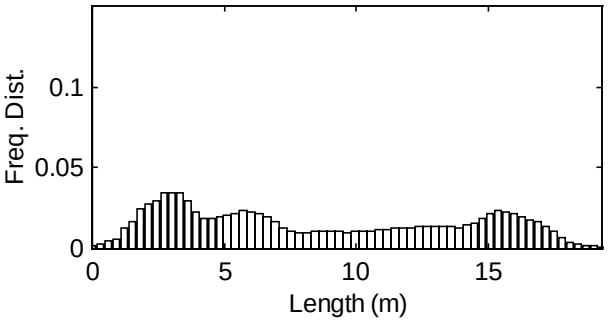

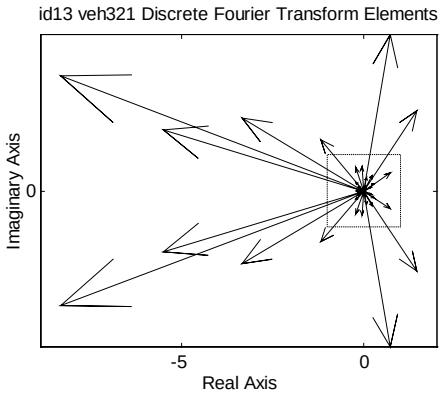

id14 veh380 ar6586 am8.10 an0.62 ev0.49 var10.94 sv-4.03 sc-0.11 mph69.98 ft37.26 mag812 l2

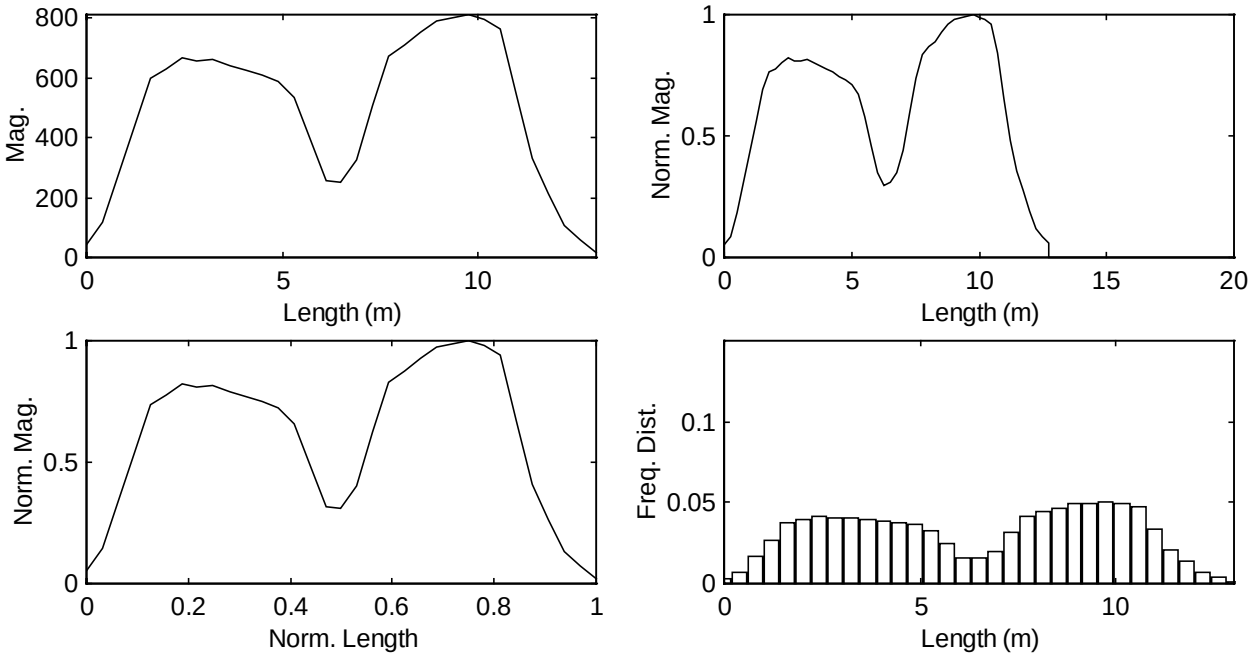

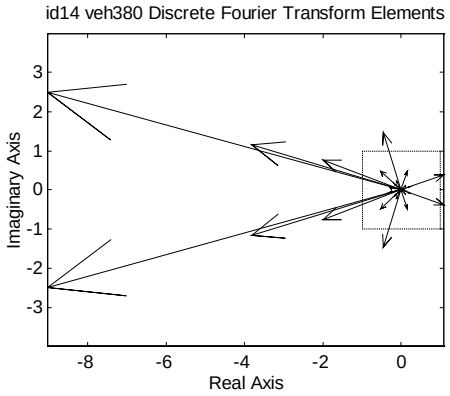

id15 veh2096 ar5237 am8.06 an0.62 ev0.49 var9.49 sv2.24 sc0.08 mph62.32 ft37.03 mag650 l3

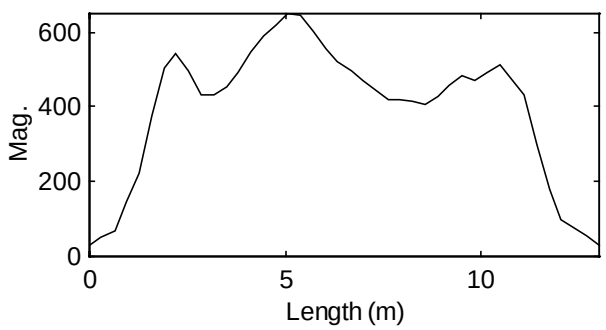

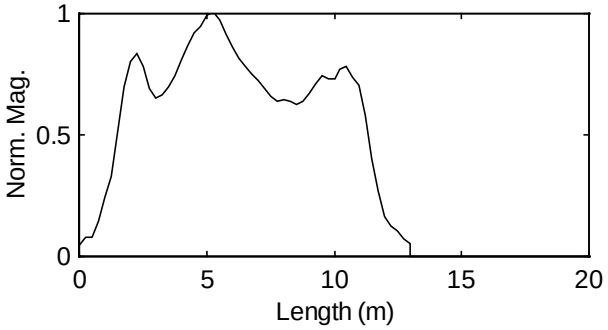

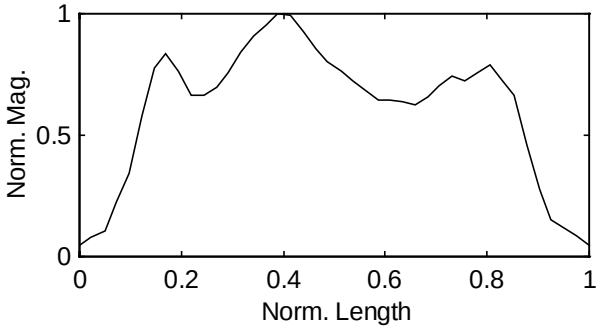

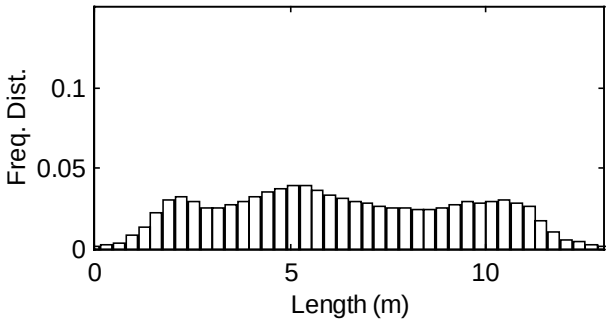

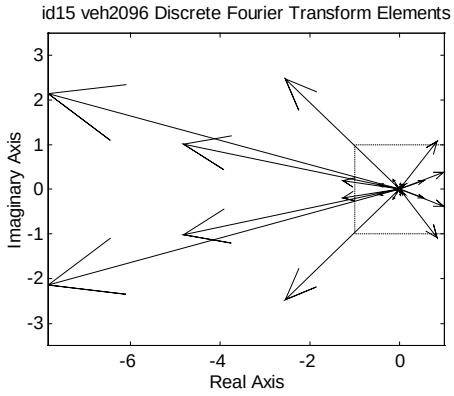

Table 1. First 14 Sorted Magnitudes of Discrete Fourier Transform Coefficients

Id	Discrete Fourier Transform Coefficient Magnitudes													
1	5.76	5.76	5.49	5.49	1.26	1.26	1.08	1.08	0.62	0.62	0.54	0.54	0.45	0.45
2	6.39	6.39	0.57	0.57	0.26	0.26	0.23	0.23	0.19	0.19	0.17	0.17	0.13	0.13
3	6.64	6.64	0.57	0.57	0.52	0.52	0.13	0.13	0.07	0.07	0.03	0.03	0.01	0.01
4	8.96	8.96	2.48	2.48	0.72	0.72	0.30	0.30	0.15	0.15	0.12	0.12	0.12	0.12
5	6.66	6.66	0.52	0.52	0.20	0.20	0.14	0.14	0.13	0.13	0.11	0.11	0.04	0.04
6	6.51	6.51	0.28	0.28	0.26	0.26	0.26	0.26	0.25	0.25	0.17	0.17	0.09	0.09
7	8.23	8.23	2.30	2.30	1.72	1.72	1.08	1.08	0.81	0.81	0.33	0.33	0.24	0.24
8	8.73	8.73	7.08	7.08	6.15	6.15	3.41	3.41	2.13	2.13	1.87	1.87	0.67	0.67
9	8.36	8.36	1.96	1.96	0.61	0.61	0.35	0.35	0.14	0.14	0.11	0.11	0.11	0.11
10	7.40	7.40	0.48	0.48	0.43	0.43	0.31	0.31	0.16	0.16	0.12	0.12	0.11	0.11
11	5.94	5.94	5.42	5.42	5.00	5.00	1.51	1.51	1.43	1.43	1.15	1.15	0.75	0.75
12	5.98	5.98	0.38	0.38	0.28	0.28	0.14	0.14	0.08	0.08	0.05	0.05	0.03	0.03
13	8.87	8.87	5.74	5.74	4.34	4.34	3.89	3.89	2.62	2.62	1.82	1.82	0.87	0.87
14	9.37	9.37	3.98	3.98	2.14	2.14	1.52	1.52	1.16	1.16	0.72	0.72	0.50	0.50
15	8.13	8.13	4.92	4.92	3.55	3.55	1.36	1.36	1.29	1.29	1.05	1.05	0.56	0.56

Table 2. First 12 Sorted Angles of Discrete Fourier Transform Coefficients

Id	Discrete Fourier Transform Coefficient Angles											
1	-2.51	2.51	-2.91	2.91	2.14	-2.14	2.06	-2.06	-1.22	1.22	-1.34	1.34
2	-3.05	3.05	-0.82	0.82	2.02	-2.02	1.78	-1.78	-1.96	1.96	0.51	-0.51
3	-2.77	2.77	-0.54	0.54	2.28	-2.28	1.90	-1.90	0.14	-0.14	1.64	-1.64
4	3.01	-3.01	-3.07	3.07	2.93	-2.93	-2.11	2.11	1.23	-1.23	2.78	-2.78
5	-2.77	2.77	-2.67	2.67	-1.07	1.07	1.87	-1.87	2.06	-2.06	-0.87	0.87
6	-3.02	3.02	-2.14	2.14	-0.43	0.43	1.60	-1.60	0.46	-0.46	2.59	-2.59
7	2.66	-2.66	2.94	-2.94	-2.04	2.04	2.15	-2.15	-0.53	0.53	-0.84	0.84
8	-1.46	1.46	2.91	-2.91	-2.81	2.81	-2.13	2.13	1.22	-1.22	-0.12	0.12
9	2.79	-2.79	-2.70	2.70	-1.42	1.42	0.92	-0.92	1.36	-1.36	-0.84	0.84
10	-2.80	2.80	0.80	-0.80	-2.46	2.46	1.19	-1.19	-0.41	0.41	1.46	-1.46
11	1.79	-1.79	-2.94	2.94	-3.13	3.13	-3.09	3.09	1.28	-1.28	-0.26	0.26
12	-2.90	2.90	-2.26	2.26	0.38	-0.38	-2.58	2.58	2.77	-2.77	-0.24	0.24
13	2.78	-2.78	2.85	-2.85	1.40	-1.40	2.60	-2.60	-0.99	0.99	-2.27	2.27
14	-2.87	2.87	2.85	-2.85	-2.78	2.78	-1.87	1.87	0.34	-0.34	-2.41	2.41
15	-2.88	2.88	2.93	-2.93	-2.38	2.38	0.91	-0.91	2.99	-2.99	0.37	-0.37

Table 3. Karhunen-Loeve Eigen values (first 9 components)

ID	1st	2nd	3rd	4th	5th	6th	7th	8th	9th
1	8.92	7.87	6.10	5.54	2.51	2.24	1.78	1.36	1.33
2	7.92	6.50	2.37	2.27	1.51	1.28	1.18	0.89	0.64
3	7.73	6.98	2.47	2.03	1.48	1.29	1.09	0.86	0.69
4	10.37	9.24	3.72	3.46	2.25	1.97	1.52	1.32	1.25
5	7.87	7.29	2.60	2.24	1.51	1.33	1.19	0.89	0.67
6	8.23	6.76	2.46	2.26	1.50	1.37	1.18	0.90	0.69
7	11.11	10.72	4.14	3.79	2.38	2.01	1.63	1.37	1.22
8	13.23	10.65	8.67	8.22	7.01	5.47	4.66	4.16	3.32
9	9.18	8.84	3.43	3.08	1.91	1.49	1.30	1.18	0.93
10	8.59	7.79	2.78	2.19	1.55	1.38	1.19	0.92	0.74
11	9.25	8.79	7.50	5.20	4.24	3.58	2.85	2.18	1.87
12	7.48	6.47	2.48	2.06	1.48	1.25	1.05	0.82	0.65
13	13.77	11.93	8.50	8.15	6.05	5.12	3.17	2.95	2.49
14	12.73	11.97	4.56	4.28	2.95	2.37	1.73	1.67	1.61
15	9.08	8.73	6.56	5.97	3.96	3.61	2.41	2.37	1.81

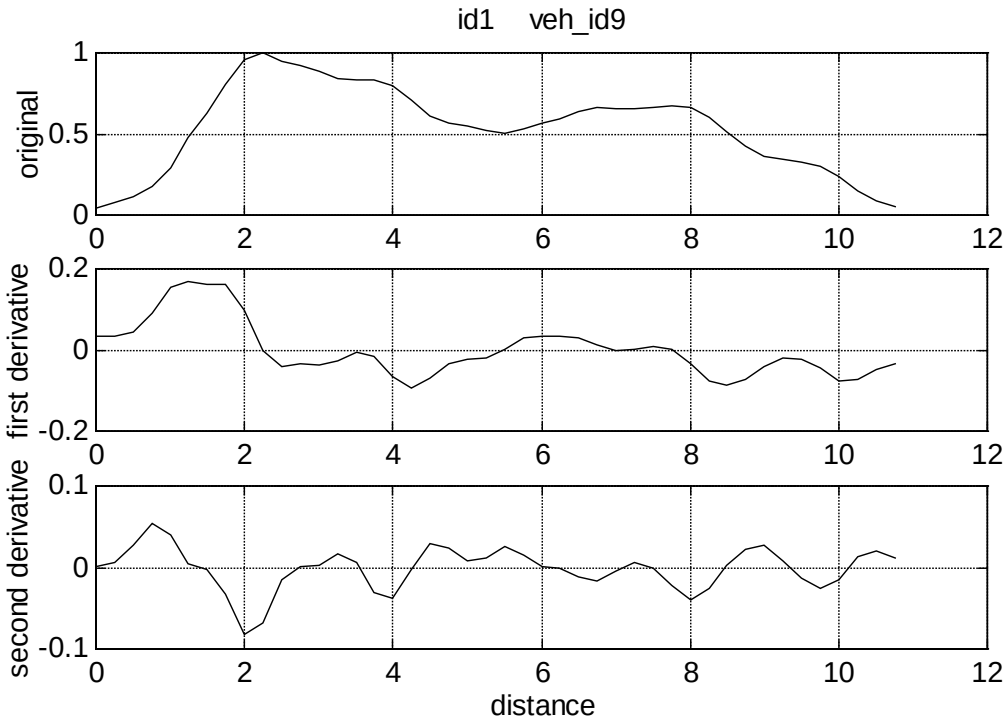
id1 veh_id9
original
first derivative
second derivative
distance

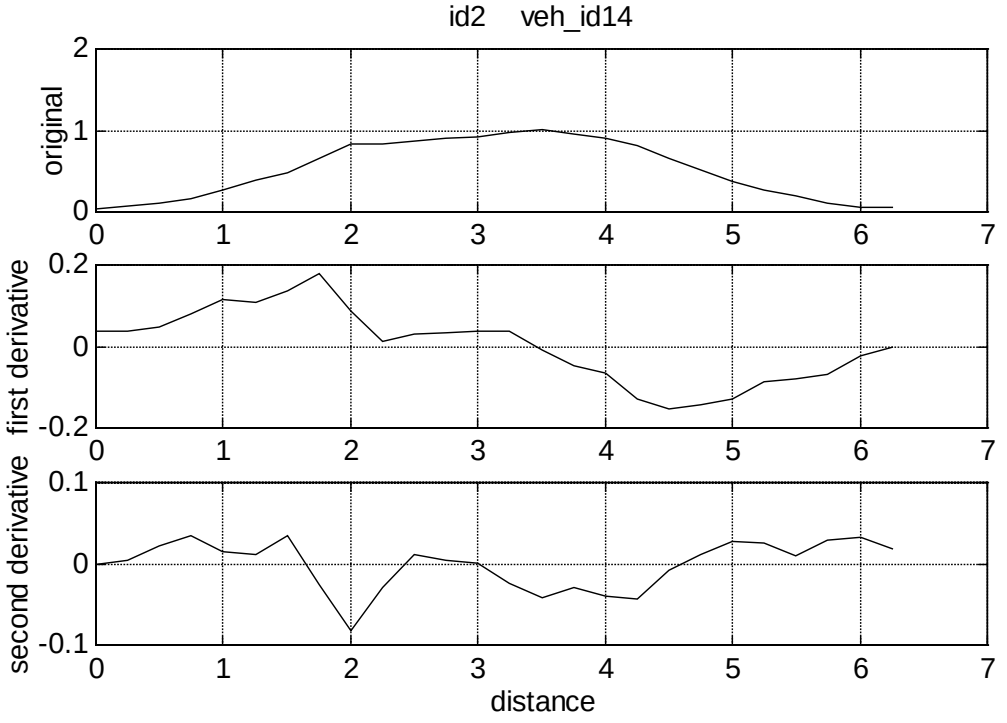
id2 veh_id14
original
first derivative
second derivative
distance

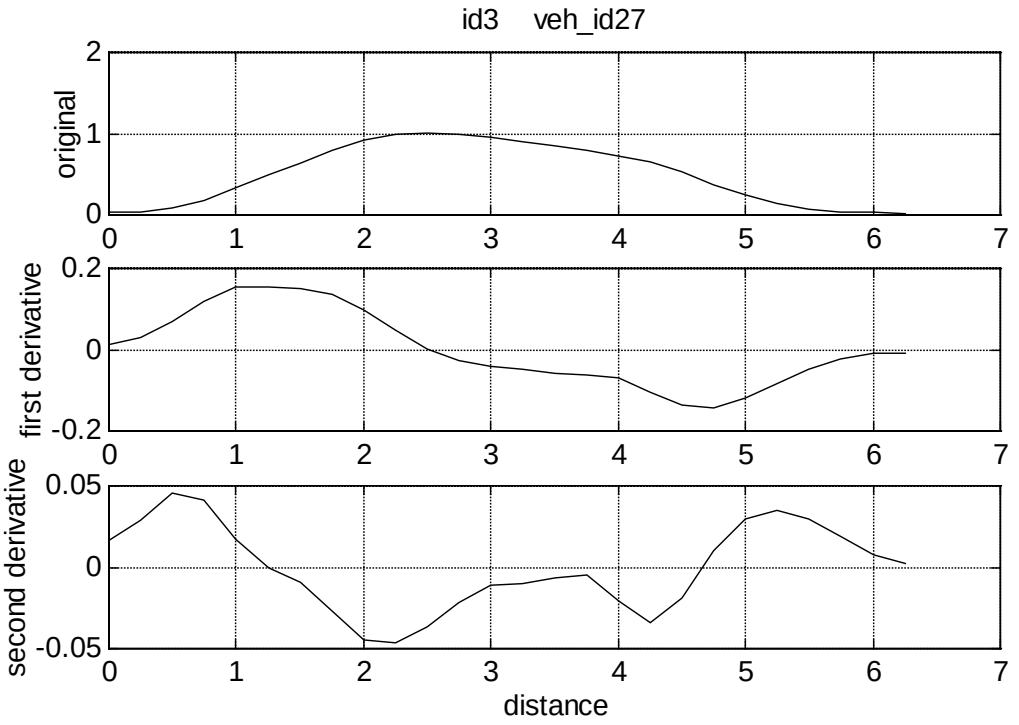
id3 veh_id27
original
first derivative
second derivative
distance

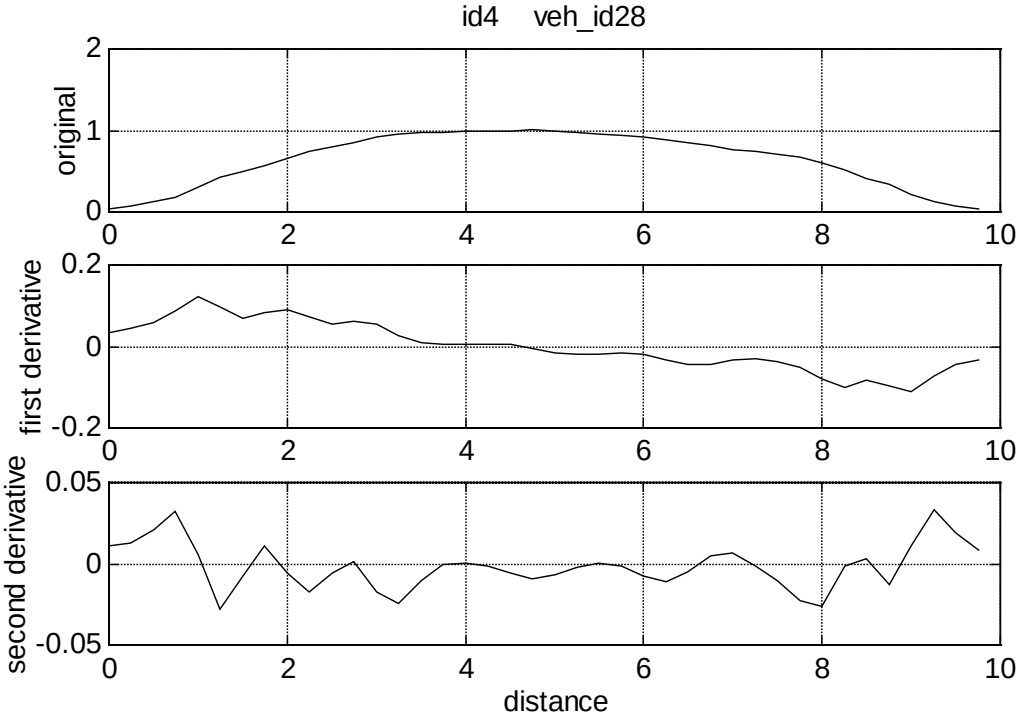
id4 veh_id28
original
first derivative
second derivative
distance

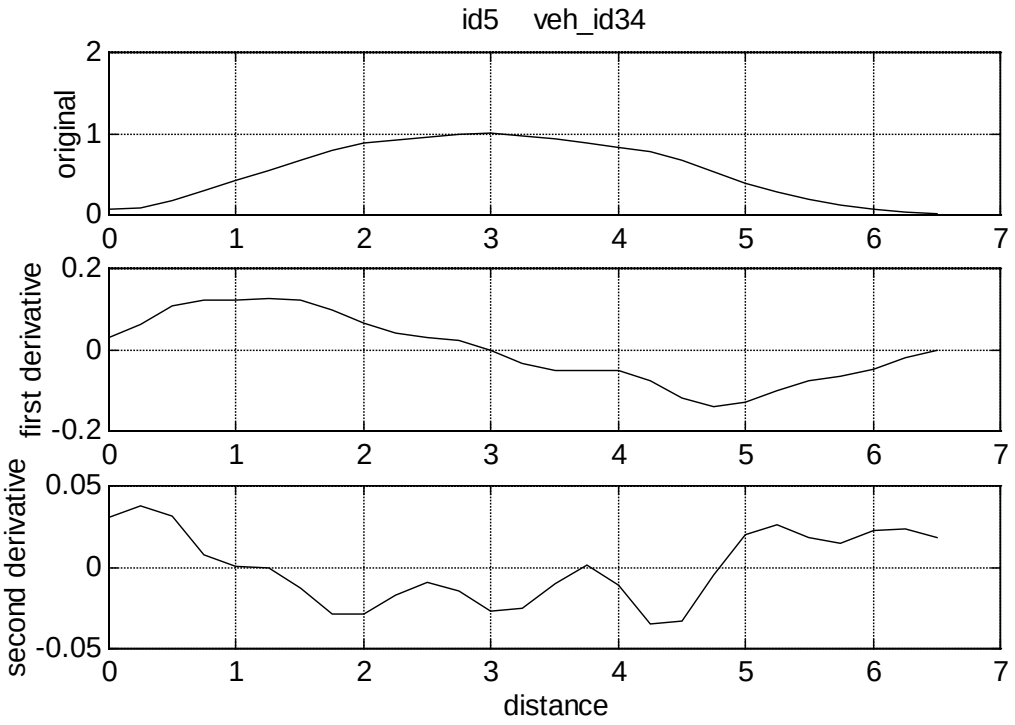
id5 veh_id34
original
first derivative
second derivative
distance

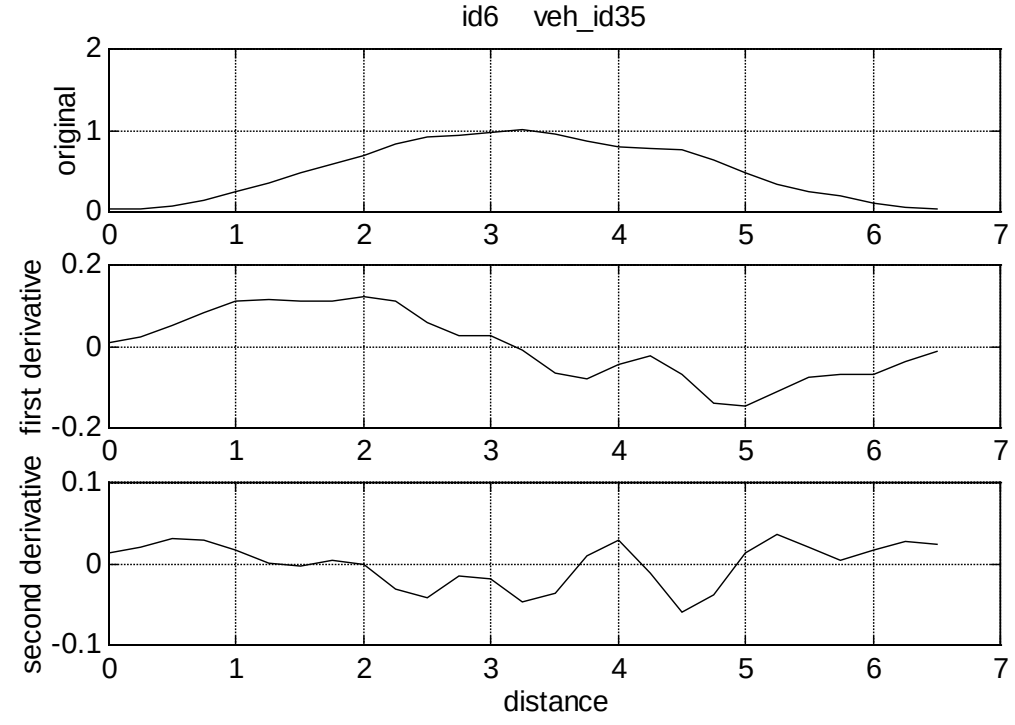
id6 veh_id35
original
first derivative
second derivative
distance

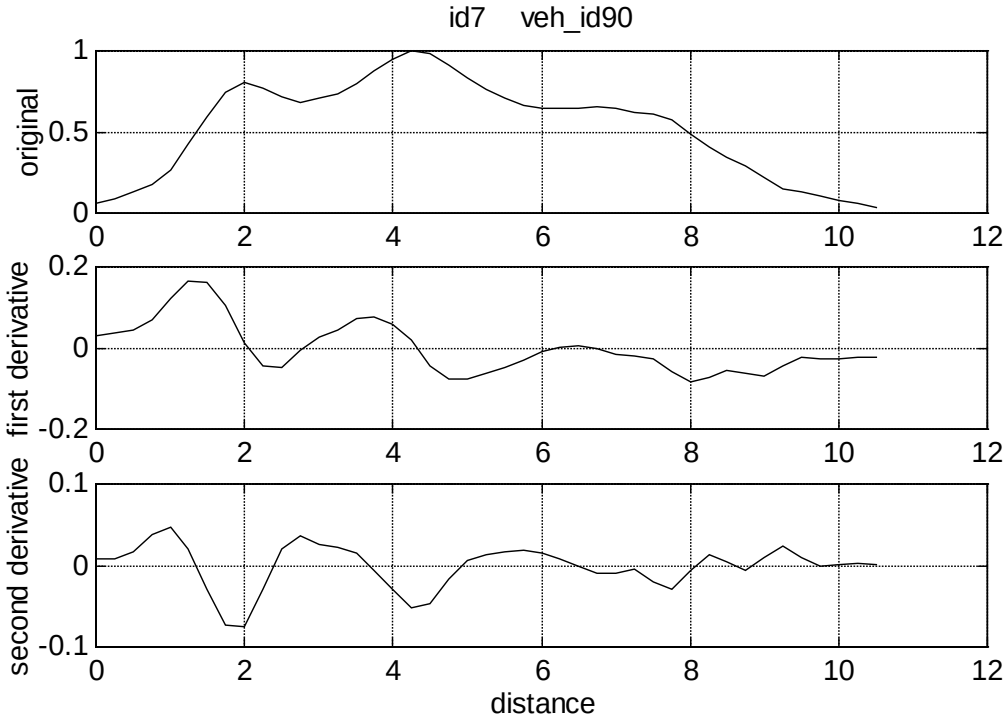
id7 veh_id90
original
first derivative
second derivative
distance

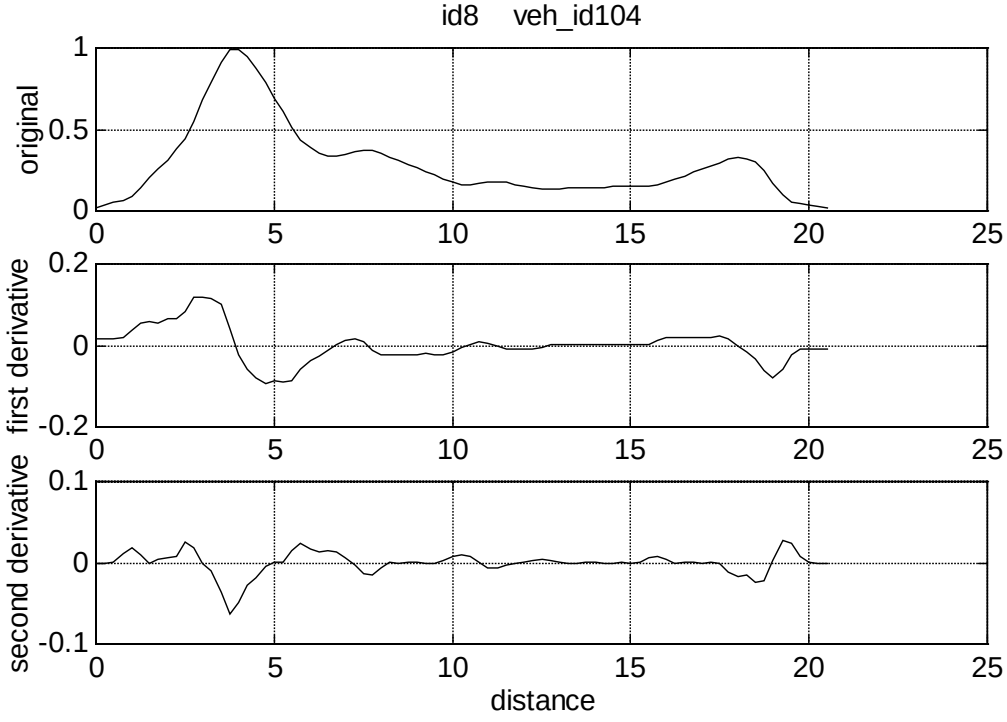
id8 veh_id104
original
first derivative
second derivative
distance

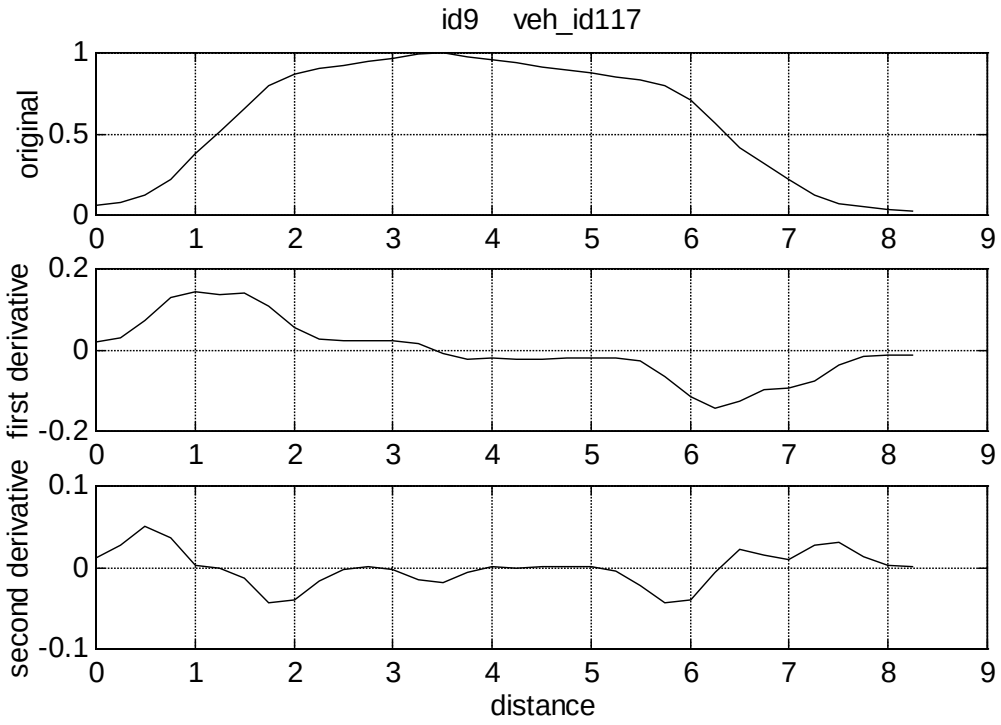
id9 veh_id117
original
first derivative
second derivative
distance

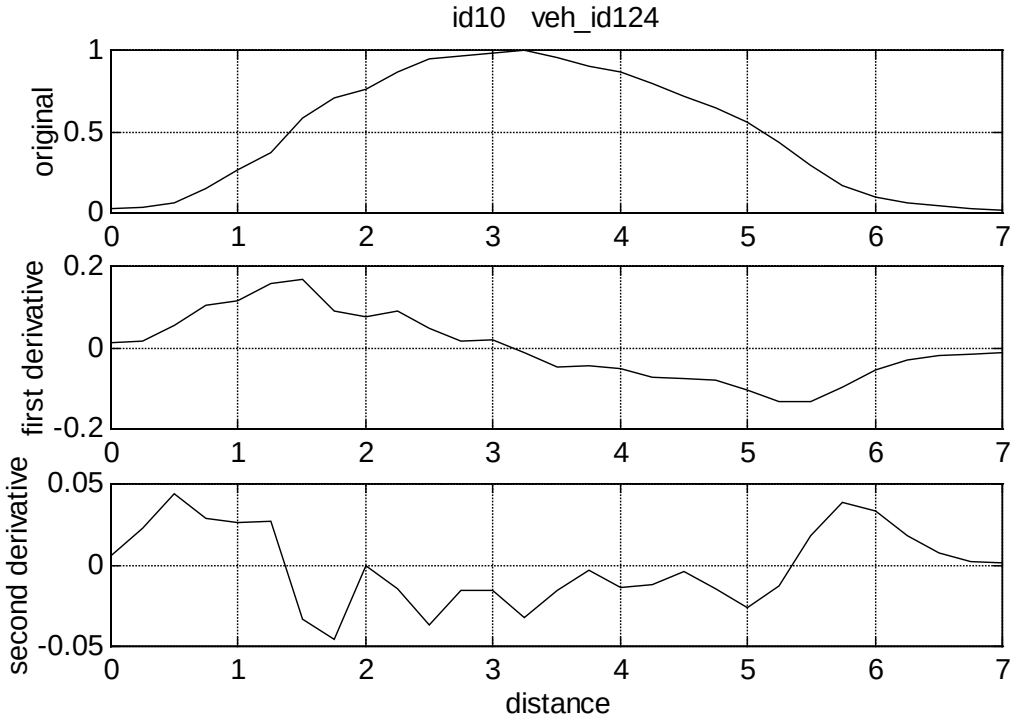
id10 veh_id124
original
first derivative
second derivative
distance

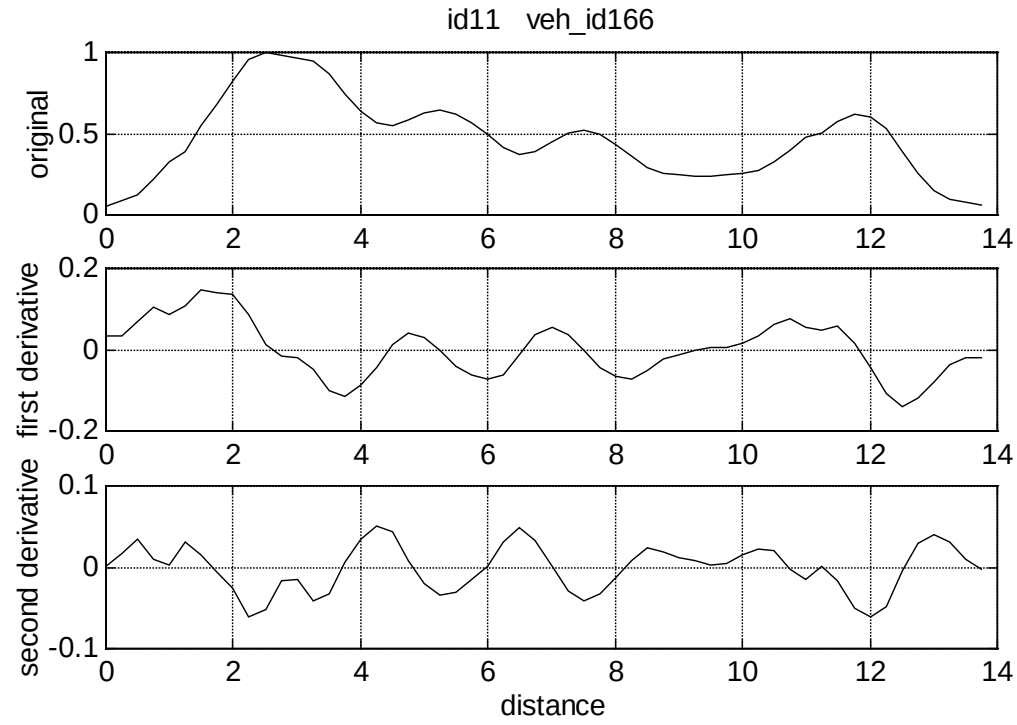
id11 veh_id166
original
first derivative
second derivative
distance

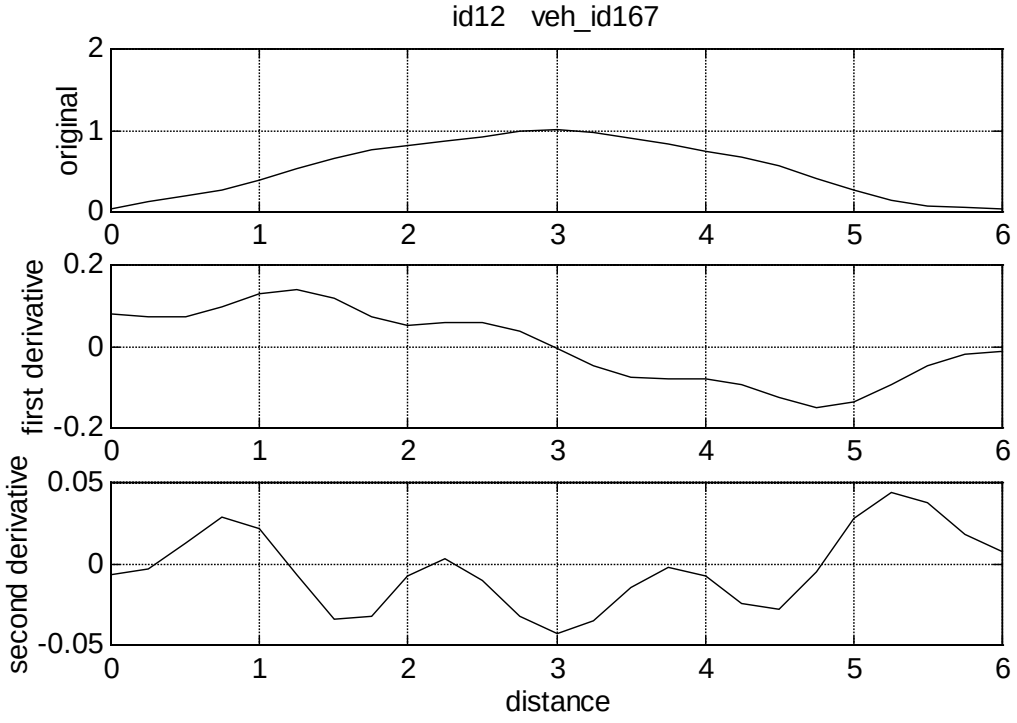
id12 veh_id167
original
first derivative
second derivative
distance

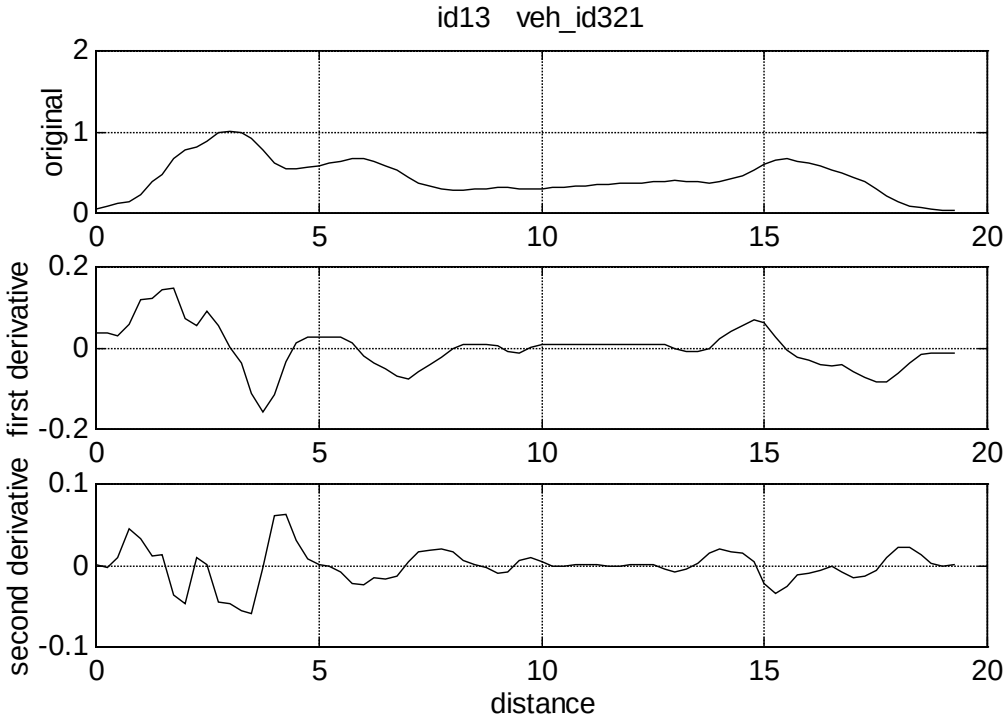
id13 veh_id321
original
first derivative
second derivative
distance

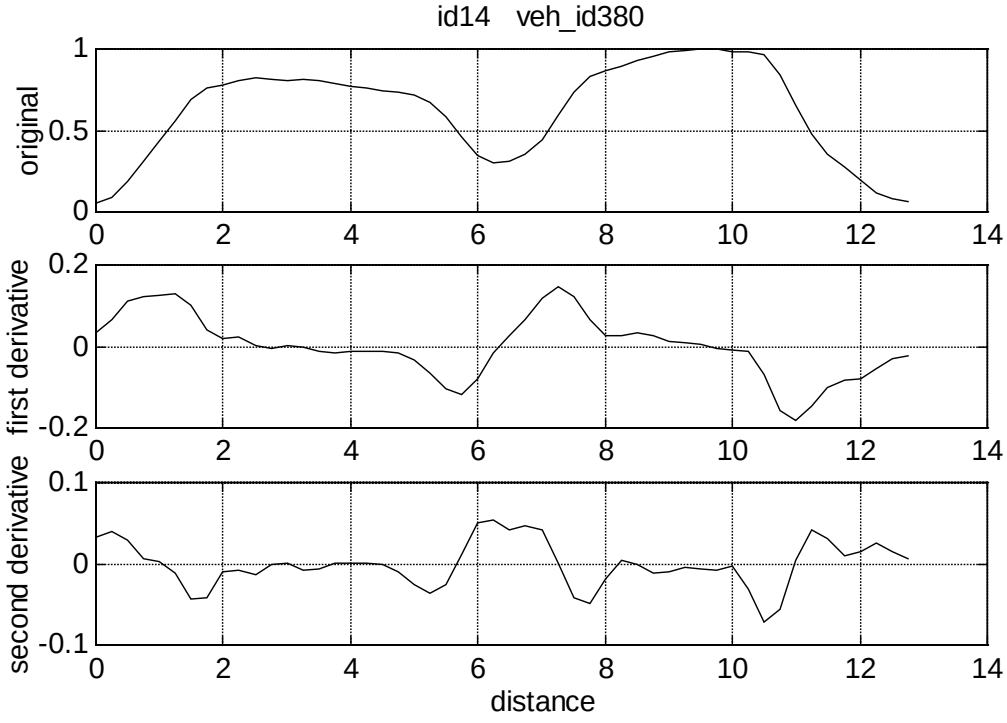
id14 veh_id380
original
first derivative
second derivative
distance

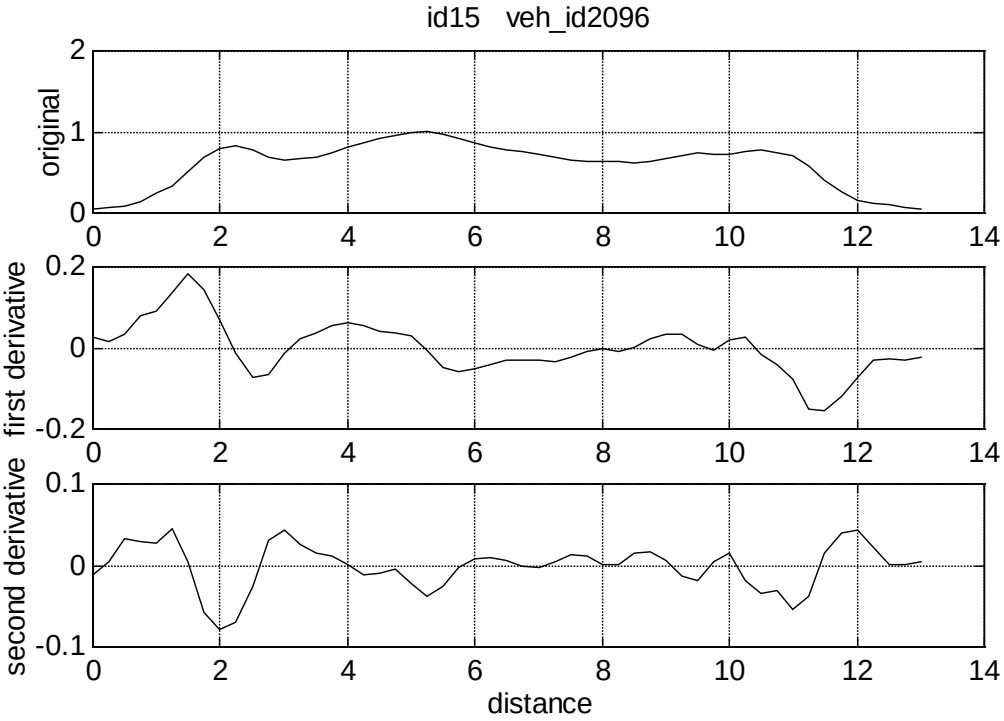
id15 veh_id2096
original
first derivative
second derivative
distance
2
1
0
0.2
0
-0.2
0.1
0
-0.1
0
2
4
6
8
10
12
14

REFERENCES

Abbas, H.M. and Fahmy, M.M. Neural Model for Karhunen-Loeve Transform with Application to Adaptive Image Compression. IEE Proceedings-I, Vol. 140, No. 2, April. Pp. 135-143. 1993.

Arendonk, J. A Comparison of Real-Time Freeway Speed Estimation Using Loop Detectors and AVI Technologies., Southwest Region University Transportation Center. Compendium: Graduate Student Papers on Advanced Surface Transportation Systems. Texas Transportation Institute. Texas A&M, College Station. Pg. J1-J40.

1996.

Athol, P. Interdependence of Certain Operational Characteristics Within a Moving Traffic Stream. Highway Research Record 72. Pg. 58-87. HRB. National Research Council. Washington, D.C. 1965.

Böhnke, P. and Pfannerstill, E. A System for the Automatic Detection of Traffic Situations. ITE Journal. Vol. 56. 1986.

Bollen, K.A. Structural Equations with Latent Variables. John Wiley & Sons. New York. Pg. 418-420. 1989.

Bow, S.T. Pattern Recognition and Image Preprocessing. Marcel Dekker, Inc. New York. 1992. Pps. 1-142.

Budget Commercial Fleet. http://www.drivebudget.com/trucks/consumer/fleet/commercial.html

Chatfield, C. The Analysis of Time Series. Fifth Edition. Chapman and Hall. London. 1996.

Chen, C.S. and Huo, K.S. Karhunen-Loeve Method for Data Compression and Speech Synthesis. IEE Proceedings-I, Vol. 138, No. 5, October. Pp. 377-380.1991.

Churchill, R.V. and Brown, J.W. Complex Variables and Applications. McGraw-Hill. New York. 1984.

Demuth, H. and Beale, M. Matlab Neural Network Toolbox User's Guide. Version 2. The Mathworks. 1997.

Fishman, G.S. Spectral Methods in Economics. Harvard University Press. Massachusetts. 1969.

Garrott, W.R., Howe, G.J., and Forkenbrock, G. An Experimental Examination of Selected Maneuvers That May Induce On-Road Untripped, Light Vehicle Rollover. National Highway Traffic Safety Administration. Washington, D.C. July. 1999.

Hall, F.L., and Persaud, B.N. Evaluation of Speed Estimation Made with Single-Detector Data from Freeway Traffic Management Systems.Transportation Research Record 1232. Pg. 9-16. 1989.

Hecht-Nielsen, R. Neurocomputing. Addison-Wesley. Massachusetts. Pps. 138-146. 1990.

ITE. Traffic Detector Handbook. 2nd Edition. Institute of Transportation Engineers. Washington, D.C. 1990.

Jacobson, L.N., Nihan, N.L., and Bender, J.D. Detecting Erroneous Loop Detector Data in a Freeway Traffic Management System. Preprint. Transportation Research Board 69th Annual Meeting. January 7-11. Washington, D.C. 1990.

Johnson, R.A. and Wichern, D.W. Applied Multivariate Statistical Analysis. Third Edition. Prentice Hall. New Jersey. 1992.

Kreizig, E. Advanced Engineering Mathematics. John Wiley & Sons. New York. Pg. 1176-1177. 1993.

Kühne, R. D. Freeway Control Using a Dynamic Traffic Flow Model and Vehicle Reidentfication Techniques. Transportation Research Record 1320. 1991. Pp. 251-259.

Kühne, R. D and Immes, S. Freeway Control Systems for Using Section-Related Traffic Variable Detection. Pacific Rim TransTech Conference, July 25-28, Seattle, Washington. 1993.

Leonard, J.D. Analysis of Large Truck Crashes on Freeway-to-Freeway Connectors. Ph.D. Dissertation. University of California, Irvine. Pp. 112-140.1991.

Lippman, R.P. An Introduction to Computing with Neural Nets. IEEE Acoustics, Speech, and Signal Processing Magazine, April 1987. Pps. 4-22. 1987.

Lu, Y. Vehicle Classification Using Infrared Image Analysis. ASCE Journal of Transportation Engineering. Vol. 118. No. 2. March/April. 1989. Pp. 223-240.

Luenberger, D.G. Linear and Nonlinear Programming. Addison-Wesley. Reading. 1989. Pps. 198-200.

Malki, H.A. and Moghaddamjoo, A. Using the Karhunen-Le've Transformation in Back-Propagation Training Algorithm. IEEE Transactions on Neural Networks, Vol. 2, No. 1, January. Pp. 162-165. 1991.

Mikhalkin, B., Payne, H.J., and Isaksen, L. Estimation of Speed from Presence Detectors. Highway Research Record 388. Pg. 73-83. HRB. National Research Council. Washington, D.C. 1972.

Nooralahiyan, A.Y. et al. A Field Trial of Acoustic Signature Analysis for Vehicle Classification. Transportation Research-C. Vol. 5. No.3/4. 1997. Pp. 165-177.

Numerical Recipes Software. Numerical Recipes in Fortran 77: The Art of Scientific Computing. Cambridge University Press. 1992.

Oppenheim, A.V. and Willsky, A.S. Signals and Systems. Prentice-Hall. New Jersey. Pps. 145-146. 1983.

PATH (Partners for Advanced Transit and Highways) Workshop on Research, Development, and Testing of Traffic Surveillance Technologies. Richmond Field Station. California. 1997.

Pursula, M. and Kosonen, I. Microprocessor and PC-based Vehicle Classification Equipments Using Induction Loops. Proceedings of the Second International Conference on Road Traffic Monitoring. IEE Publication Number 299. London. February 7-9. Pg. 24-28.1989.

Pursula, M. and Pikkarainen, P. A Neural Network Approach to Vehicle Classification with Double Induction Loops. Proceedings of the 17th ARRB Conference. Part 4. 1994. Pp. 29-44.

Schalkoff, R. Pattern Recognition: Statistical, Structural, and Neural Approaches. John Wiley and Sons. New York. 1992.

Schwartz Electro-Optics, Inc. Autosense II. Product literature. http://www.seord.com/comsen/autosenseII.htm. 1998.

Steuer, R.E. Multiple Criteria Optimization: Theory, Computation, and Application. John Wiley & Sons. New York. 1986. Pps. 165-183, 196.

Sun, C., Ritchie, S., and Tsai, W. Algorithm Development for Derivation of Section-Related Measures of Traffic System Performance. Transportation Research Record 1643. 1998. Pps. 171-180.

Sun, C., Ritchie, S., Tsai, K., and Jayakrishnan, R. Use of Vehicle Signature Analysis and Lexicographic Optimization for Vehicle Reidentification on Freeways. Forthcoming in Transportation Research Part C. 1999.

Taylor, S.S. Inductive Loop Detector Functions. Department of Traffic, City of Los Angeles. Staff Report No. 53.08. February 4. 1972.

Transport Data Systems. Automatic Vehicle Classification System. Product literature. Rhassel@IX.netcom.com. 1998.

USDOT. Comprehensive Truck Size and Weight Study. June. http://www.fhwa.dot.gov/reports/tswstudy/TSWfinal.htm 1997

Wannacott, T.H. and Wannacott, R.J. Introductory Statistics. Wiley & Sons. New York. 1990.

Wardrop, J.G. Some Theoretical Aspects of Road Traffic Research." Proceedings Instn. Div. Engrs., Vol. 1, No. 2. Pg. 325. 1952.

Wasserman, P.D. Neural Computing: Theory and Practice. Van Nostrand Reinhold. New York. Pps. 220-222. 1989.

Wei, C. et al. Vehicle Classification Using Advanced Technologies. Transportation Research Record 1551. November. 1996. Pp. 45-50.

White, H. A Heteroscedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity." Econometrica, 48, pp. 817-838. 1980.

Wilshire, R., Black, R., Grochoske R., and Higinbotham, J. Traffic Control Systems Handbook. Institute of Transportation Engineers. Washington, D.C. Report No. FHWA-IP-85-12. April. 1985.

Woods, D.L., Cronin, B.P., and Hamm, R.A. Speed Measurement with Inductance Loop Speed Traps. Texas Transportation Institute. Research Report FHWA/TX-95/1392-8. Texas A&M. College Station. 1994.

Yuan, X. et al. Computer Vision System for Automatic Vehicle Classification. ASCE Journal of Transportation Engineering. Vol. 120. No. 6. November/December. 1994. Pp. 861-876.